



Fig. 1

1. TAP Cross-Server Queries

Markus Demleitner
msdemlei@ari.uni-heidelberg.de

- A word on TAP
- Cross-server queries and a bottleneck
- The idea of the shortcut...
- ... and the practice with STILTS and pyVO.

(cf. Fig. 1)

2. A Word on TAP

TAP = Table Access Protocol: Rules for sending database queries to remote servers, monitor the execution, and fetch the results.

Two modes:

- synchronous (sync) – send your query and keep a connection to the server until the results come in.
- asynchronous (async) – send your query and get a URL at which you can watch its progress and result at your leisure.

async is more than just "use it when your queries take a while"!

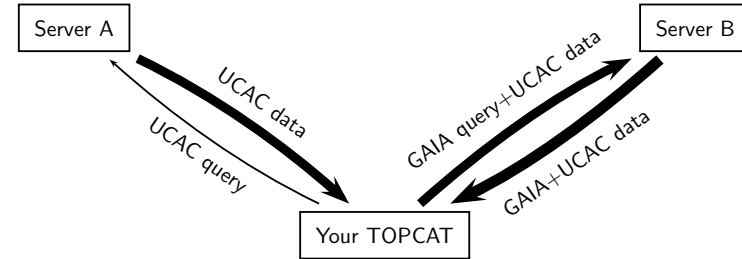


Fig. 2

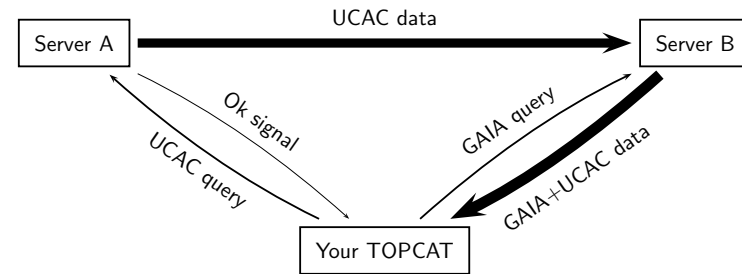


Fig. 3

3. Cross-Server Queries: Plain

Imagine you're matching something from UCAC4 on server A with Gaia DR 2 on server B. Here's the conventional way:

(cf. Fig. 2)

You're transferring the UCAC data twice for no good reason. Wouldn't be be smarter to...

4. Cross-Server Queries: Smart

... transfer the UCAC data directly from server A to server B, like this:

(cf. Fig. 3)

Save an extra download and upload. Also, server A and server B may have a thick pipe between them.

5. Is that possible?

Sure. In TAP from day 1.

Idea: Run an async query on server A and use the result URL as an upload on server B.

TOPCAT doesn't support that, though. But pyVO does. And STILTS almost.

6. Cross-Server Uploads in pyVO

Sketch:

```
svc1 = TAPService(url1)
svc1.submit_query("SELECT FROM ucac4")
result = svc1.run()
```

```
svc2 = TAPService(url2)
svc2.submit_query("SELECT FROM gaia JOIN upload",
    upload=result.url)
result = svc2.run()
samp.send_table_to(result, "topcat")
```

In practice, there's a bit more technics around it, but not much. See attachment.

See PDF attachment(s): [withpyvo.py](#)

7. Cross-Server Uploads in STILTS

STILTS can upload from a URL. In async TAP, the job URL always is <job url>/results/result.

Sketch:

```
JOBURL=$(stilts tapquery tapurl=UCAC_URL omode=joburl UCAC-QUERY)
```

```
stilts tapquery omode=topcat \
    tapurl=GAIA_URL \
    nupload1 upload1=$JOBURL/results/result \
    GAIA-QUERY
```

In reality, unfortunately, it doesn't work like this because STILTS so far insists on retrieving the uploads and reformatting them.

See the attached [withstilts.sh](#) for a rather whacky workaround involving a bit of shell. Be sure to let STILTS's author, Mark Taylor, know when you'd like to do this without magic.

See attached file for how you can do this already in real life.

See PDF attachment(s): [withstilts.sh](#)

8. Open Issue

As usual, things get ugly when access restrictions are in effect: If server A requires authentication, server B needs to get the credentials.

In 2010, the IVOA has published the Credential Delegation Protocol CDP that uses X.509 proxy certificates to solve this.

I'm not aware of living implementations of this.

OAuth lets you do similar things. I don't think any servers implement that, either.

If you're running authenticated TAP services: How about prototyping it?

9. More Info

Slides and scripts for this talk:

<http://docs.g-vo.org/talks/2019-ag-tap.pdf>

A course on pyVO: <http://docs.g-vo.org/pyvo>