

CDS xmatch service updates

Thomas Boch
François-Xavier Pineau
Sébastien Derrière



*IVOA Interop Meeting@Naples
May 17 2011 - Apps 2 - T. Boch et al.*



Service presentation

- Purpose
 - **Positional cross-identification** of sources in 2 tables among:
 - VizieR tables (including very large surveys)
 - SIMBAD
 - User-uploaded tables
 - Focus on **large-scale** cross-match
 - UI : web application on top of UWS service

Demo

CDS X-Match Service [Tables management](#) [X-match](#) thomas [Preferences](#) [Admin](#) [Logout](#)

Choose tables to cross-match

NOMAD USNOB1

[My store](#) [VizieR](#) [SIMBAD](#) [My store](#) [VizieR](#) [SIMBAD](#)

Cross-match criteria

By position
Radius:

By position including error
Sigma: (completeness: 99.73 %)
Max. distance:

Cross-match area

All sky

Cone
Center:
Radius:

Healpix cell
Nside:
Index:

Visualize and manage your cross-match jobs

List of X-match jobs

Table 1	Table 2	Options	Begin	Status	Actions
DENIS	2MASS	fixed radius <input style="font-size: small; vertical-align: middle;" type="button" value="+"/>	13/05/2011 at 14:12	completed <input style="font-size: small; vertical-align: middle;" type="button" value="i"/>	<input type="button" value="Get result"/> <input type="checkbox"/>

For the selected job(s):

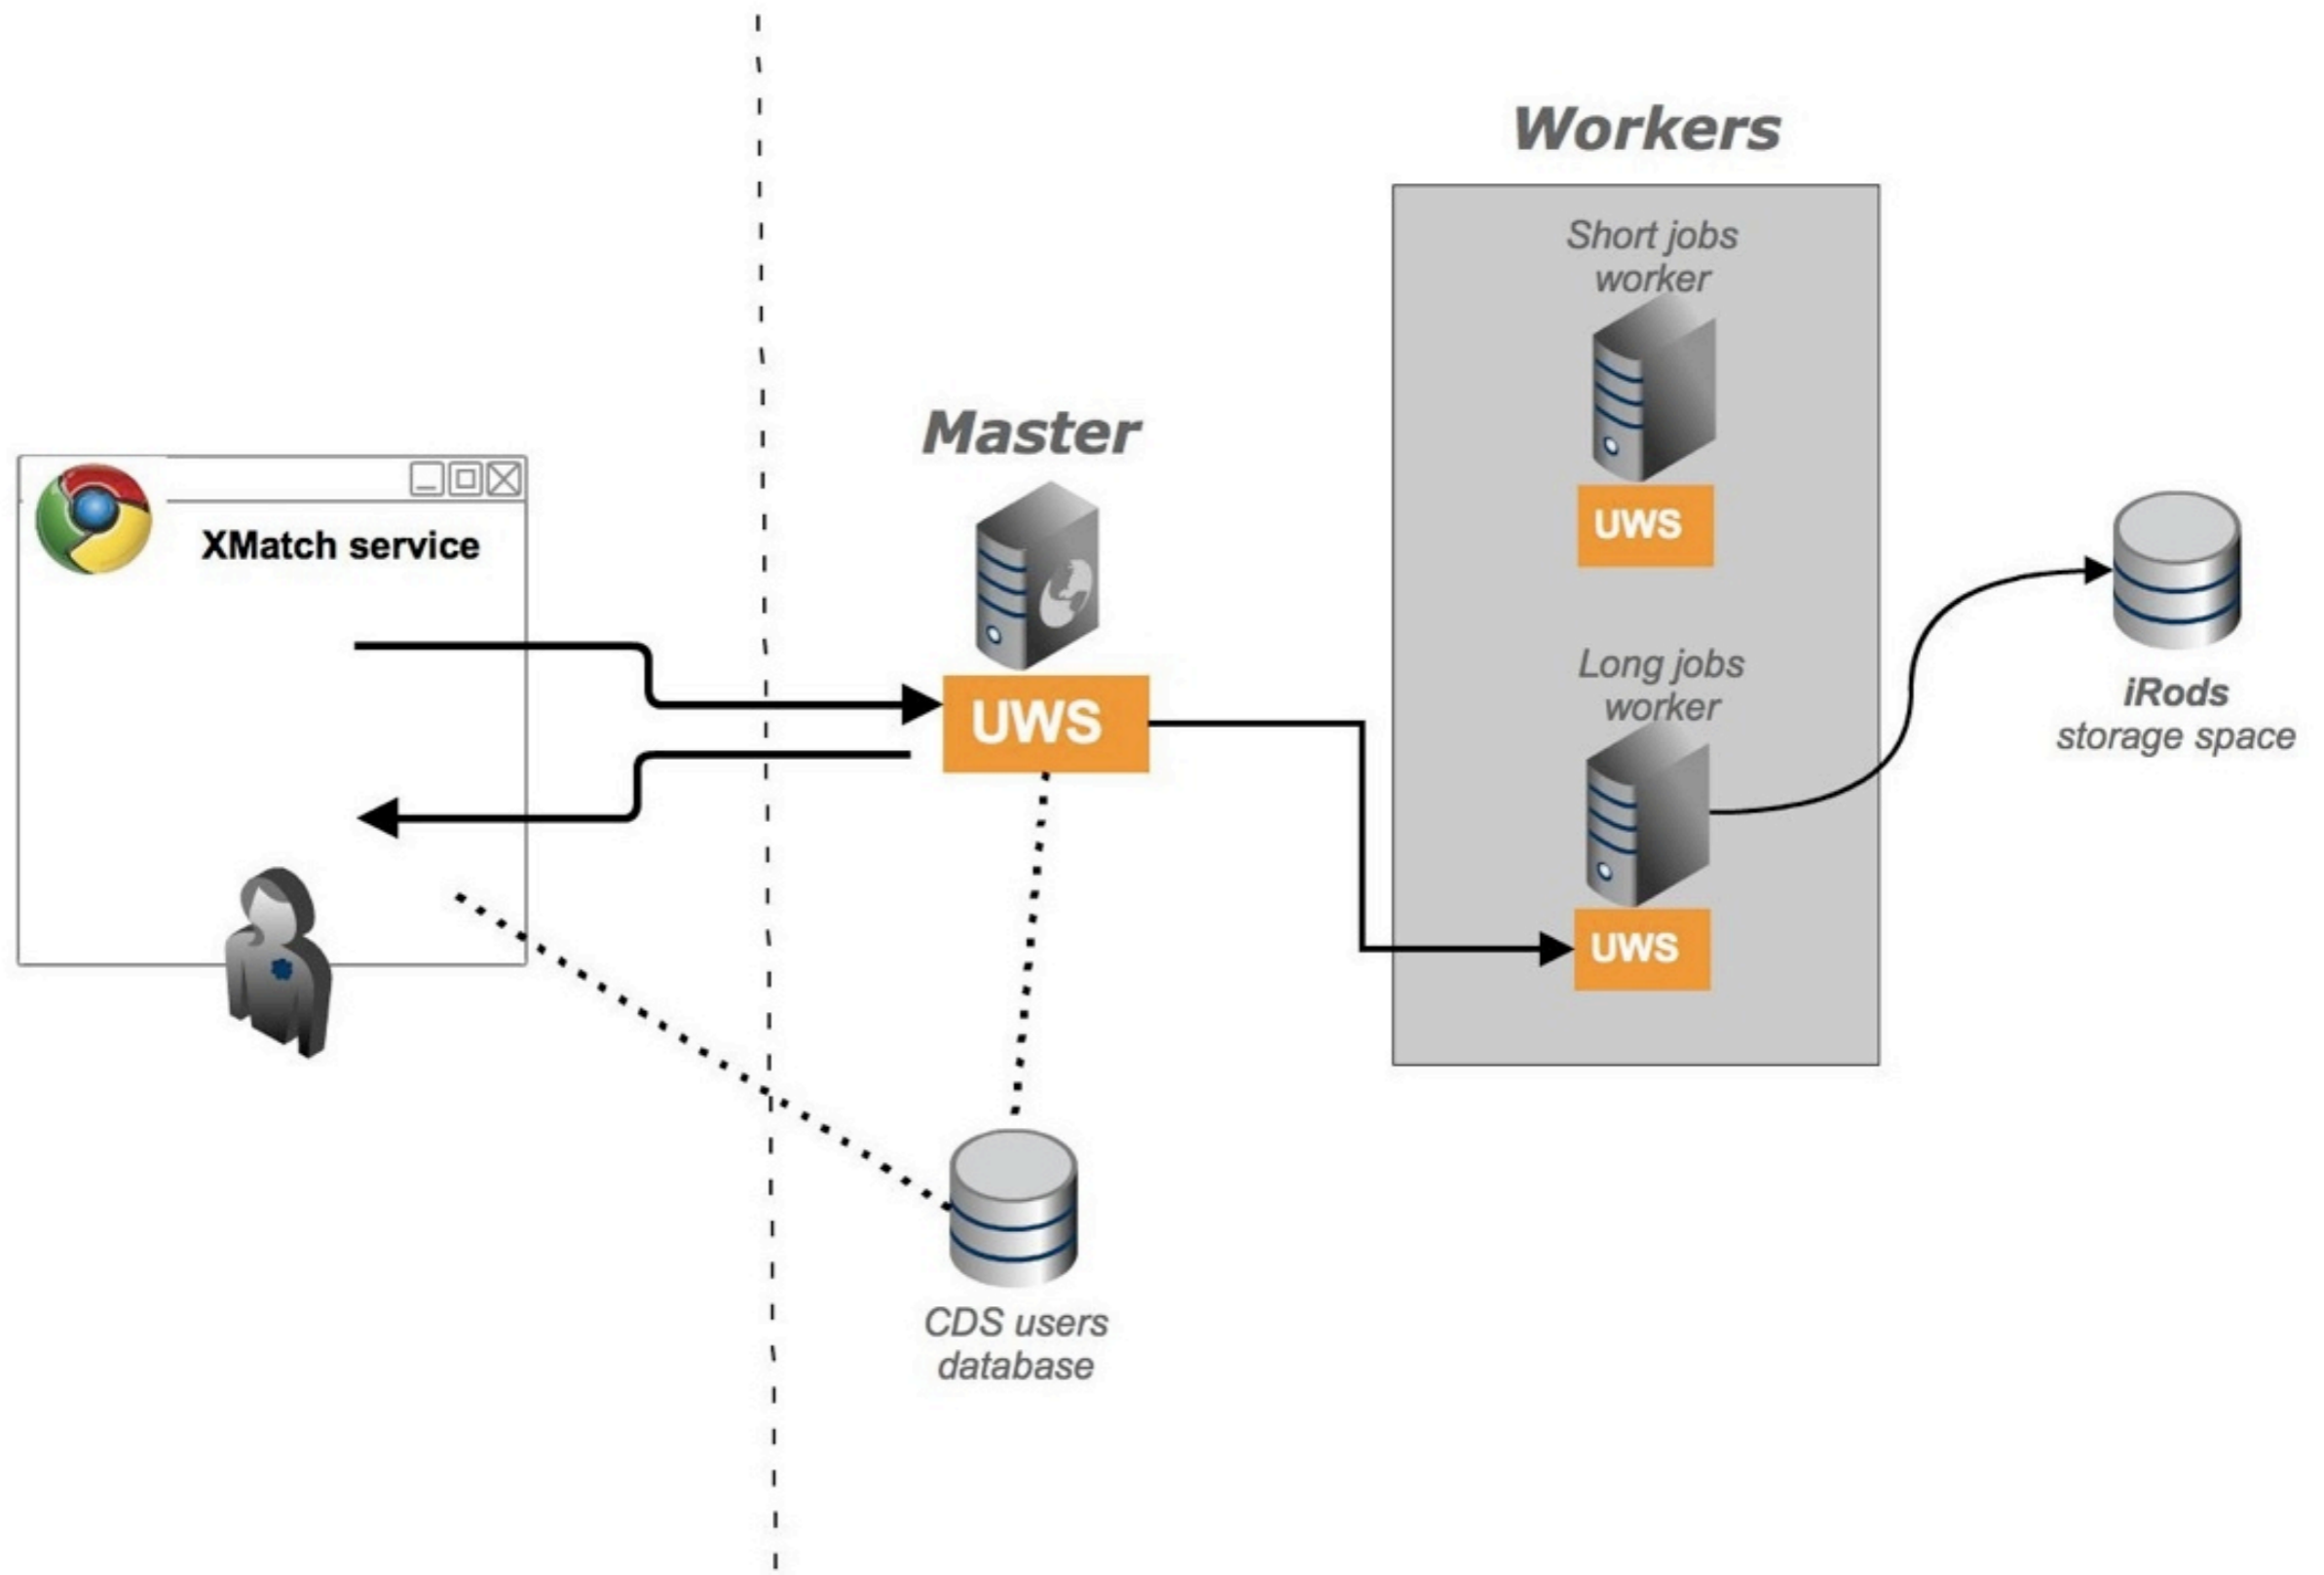
Job executed in 1h2min9s
10min59s to correlate
51min9s to generate file
Result: 272,005,494 rows



IVOA Interop@Naples
T. Boch, CDS



Architecture



Updates

- Since Nara Interop:
 - cross-match jobs **dispatched on 2 machines** (one for short jobs, one for long jobs)
 - inclusion of **VizieR tables > 10M rows** (USNOB1, 2MASS, SDSS7, NOMAD, PPMXL, GLIMPSE, ...)
 - cross-match on all-sky **or in a cone** (position or object name + radius) **or for a given HEALPix cell**
 - choice of position+errors metadata (for user-uploaded tables)
 - faster HDDs (15,000 rpm)
 - faster output generation

Service performances

Cross-match at 5 arcsec :

Table 1	Table 2	Computation time	Result generation	Result size	Total time
SDSS <i>357M rows</i>	2MASS <i>470M rows</i>	7 min	12 min	13 GB	19 min
DENIS <i>355M</i>	2MASS <i>470M</i>	11 min	51 min	58 GB	1 hour 2 min
GLIMPSE <i>104M</i>	NOMAD <i>1.1 billion</i>	6 min	17 min	19 GB	23 min
SIMBAD <i>5M</i>	USNOB1 <i>1 billion</i>	3 min	1 min	1 GB	4 min
SIMBAD <i>5M</i>	PPMX <i>18M</i>	20 seconds	20 seconds	440 MB	40 sec



IVOA Interop@Naples
T. Boch, CDS



I/O limitations

- When computing the associations :
 - **Reading input catalogues** is the main limiting factor
- When generating the result file :
 - Performance mainly limited by the **network bandwidth between workers and iRods** ($\approx 12\text{Mbyte/s}$)
- When the user downloads the result :
 - Downloading SDSS-2MASS result at 1Mbyte/s takes 3.7 hours

Hardware



- Worker 1:
 - 2x Quad Core 2.27 GHz → 16 threads
 - RAM : 24GB @ 1333MHz
- Worker 2:
 - 2x Six Core 2.27 GHz → 24 threads
 - RAM : 32GB @ 1333MHz
- 6TB RAID5 array with 15,000 rpm disks
 - Read : 570 MB/s
 - Write : 130 MB/s

CDS XMatch service updates: the engine

François-Xavier Pineau¹, Thomas Boch¹ and Sebastien Derrière¹

¹CDS, Observatoire Astronomique de Strasbourg

Interop Napoli, 17 May 2011




The xmatch engine : 2 major components

A catalog file format (*CatFile*)

- binary data
- compressed data
- Healpix sorted and indexed sources
- 2 formats: by rows (.rcf) , by blocks (.bcf)

A correlation engine

- Multi-threaded special kd-trees (ra, dec)
- Healpix partitioning (for large catalogs)
- *CatFile* (for large catalogs)
 - ▶ for the correlation (.bcf)
 - ▶ to build the output (.rcf)
- *STIL* to read VOTable, FITS, ...




Healpix hierarchical sky partitioning

The xmatch engine : 2 major components

A catalog file format (*CatFile*)


- binary data
- compressed data
- Healpix sorted and indexed sources
- 2 formats: by rows (.rcf) , by blocks (.bcf)



Healpix hierarchical sky partitioning




A correlation engine

- Multi-threaded special kd-trees (ra, dec)
- Healpix partitioning (for large catalogs)
- *CatFile* (for large catalogs)
 - ▶ for the correlation (.bcf)
 - ▶ to build the output (.rcf)
- *STIL* to read VOTable, FITS, ...



CatFile catalog file format

- (compressed) data stored by rows
 - ▶ as few random access as possible
 - (compressed) data stored by blocks
 - ▶ read as few bytes as possible
- data close on file / sources closed on sky
 - healpix indexed
 - ▶ \rightsquigarrow direct access to a pixel data
 - creation from .csv: $\sim 1h$ (SDSS)



XMatch engine: correlation modes

Correlation modes

- 3 catalog sizes (arbitrary)
 - small: < 150 000 sources
 - medium: < $20 \cdot 10^6$ sources
 - big: > $20 \cdot 10^6$ sources
- correlation modes
 - ▶ small vs medium, medium vs medium
 - ✦ brute-force (1 kd-tree)
 - ▶ small vs big
 - ▶ medium vs big, big vs big

Results (xmatch only)

- medium vs medium
 - ▶ SIMBAD ($5 \cdot 10^6$) vs PPMX ($18 \cdot 10^6$): 20s
- small vs big
 - ▶ Downes (1830) vs 2MASS ($470 \cdot 10^6$): 10s, <1s (disk cache)
- medium vs big
- big vs big

XMatch engine: correlation modes

Correlation modes

- 3 catalog sizes (arbitrary)
 - small: $< 150\,000$ sources
 - medium: $< 20 \cdot 10^6$ sources
 - big: $> 20 \cdot 10^6$ sources
- correlation modes
 - ▶ small vs medium, medium vs medium
 - ★ brute-force (1 kd-tree)
 - ▶ small vs big
 - compute Healpix touched cells
 - brute-force (1 kd-tree)
 - ▶ medium vs big, big vs big
 - cell by cell (162 Healpix cells)
 - compute Healpix touched cells?
 - brute-force (1 kd-tree by cell)

Results (xmatch only)

- medium vs medium
 - ▶ SIMBAD ($5 \cdot 10^6$) vs PPMX ($18 \cdot 10^6$): 20s
- small vs big
 - Downes (1830) vs 2MASS ($470 \cdot 10^6$): 10s, $< 1s$ (disk cache)
- medium vs big
 - Tycho2 ($2 \cdot 10^6$) vs 2MASS ($470 \cdot 10^6$): 2min
- big vs big
 - SDSS7 ($350 \cdot 10^6$) vs 2MASS ($470 \cdot 10^6$): 4min
 - 2MASS ($470 \cdot 10^6$) vs USNOB1 ($1 \cdot 10^8$): 20min

XMatch engine: correlation modes

Correlation modes

- 3 catalog sizes (arbitrary)
 - small: $< 150\,000$ sources
 - medium: $< 20 \cdot 10^6$ sources
 - big: $> 20 \cdot 10^6$ sources
- correlation modes
 - ▶ small vs medium, medium vs medium
 - ★ brute-force (1 kd-tree)
 - ▶ small vs big
 - ★ compute Healpix touched cells
 - ★ brute-force (1 kd-tree)
 - ▶ medium vs big, big vs big
 - cell by cell (192 Healpix cells)
 - compute Healpix touched cells?
 - brute-force (1 kd-tree by cell)

Results (xmatch only)

- medium vs medium
 - ▶ SIMBAD ($5 \cdot 10^6$) vs PPMX ($18 \cdot 10^6$): 20s
- small vs big
 - ▶ Downes (1830) vs 2MASS ($470 \cdot 10^6$): 10s, $< 1s$ (disk cache)
- medium vs big
 - Tycho2 ($2 \cdot 10^6$) vs 2MASS ($470 \cdot 10^6$): 2min
- big vs big
 - SDSS7 ($350 \cdot 10^6$) vs 2MASS ($470 \cdot 10^6$): 4min
 - 2MASS ($470 \cdot 10^6$) vs USNOB1 ($1 \cdot 10^9$): 20min

XMatch engine: correlation modes

Correlation modes

- 3 catalog sizes (arbitrary)
 - small: $< 150\,000$ sources
 - medium: $< 20 \cdot 10^6$ sources
 - big: $> 20 \cdot 10^6$ sources
- correlation modes
 - ▶ small vs medium, medium vs medium
 - ★ brute-force (1 kd-tree)
 - ▶ small vs big
 - ★ compute Healpix touched cells
 - ★ brute-force (1 kd-tree)
 - ▶ medium vs big, big vs big
 - ★ cell by cell (192 Healpix cells)
 - ★ compute Healpix touched cells?
 - ★ brute-force (1 kd-tree by cell)

Results (xmatch only)

- medium vs medium
 - ▶ SIMBAD ($5 \cdot 10^6$) vs PPMX ($18 \cdot 10^6$): 20s
- small vs big
 - ▶ Downes (1830) vs 2MASS ($470 \cdot 10^6$): 10s, $< 1s$ (disk cache)
- medium vs big
 - ▶ Tycho2 ($2 \cdot 10^6$) vs 2MASS ($470 \cdot 10^6$): 2min
- big vs big
 - ▶ SDSS7 ($350 \cdot 10^6$) vs 2MASS ($470 \cdot 10^6$): 4min
 - ▶ 2MASS ($470 \cdot 10^6$) vs USNOB1 ($1 \cdot 10^9$): 20min

What's next?

Service

- Beta testing phase in June
 - ▶ interested in testing the service: thomas@astro.unistra.com
- First release in summer

Futur developments

- Service:
 - ▶ to take into account proper motions
 - ▶ add constraints on both catalogs and the result (color, mag, ...)
- CatFile:
 - ▶ add a proper motions block
 - ▶ add constraints on catalog columns
 - ▶ allow user plugins for compression

What's next?

Service

- Beta testing phase in June
 - ▶ interested in testing the service: thomas@astro.unistra.com
- First release in summer

Futur developments

- Service:
 - ▶ to take into account proper motions
 - ▶ add constraints on both catalogs and the result (color, mag, ...)
- CatFile:
 - ▶ add a proper motions block
 - ▶ add constraints on catalog columns
 - ▶ allow user plugins for compression