PROFESSIONAL Postgres

PostgreSQL Universal Database

Oleg Bartunov

Moscow University,
Postgres Professional

Pgconf.ru, Febraury 5, 2018, Moscow, Russia

Postgres PROFESSIONAL

How to choose a right database ?

- People usually choose a database looking on
  - Functionality, Performance
  - Availability - License, price
  - Local expertise, Personal experience
  - Compatibility to existing environment
  - Support
  - Other
- People start to think about one important feature of database after the project needs
  - Need new functionality, Better performance
- Project is in production, no way to change database
  - Starting to use various ugly «solutions»
  - System works, but looks pretty strange

Your system with wrong database

PostgreSQL Universal Database

- Any project could start with PostgreSQL
- PostgreSQL is a reliable and stable database with rich functionality and long history
- PostgreSQL has liberal BSD license, cross platform (~30)
- Developed by international community, no vendor lock

- PostgreSQL is **EXTENSIBLE**, this is the very important feature, which people miss ! It allow database to support
  - New workloads
  - New functionality
  - New environment
  - Often without restarting a server, no need core programmer.

Case 1: Web

- 1996: Start using Postgres on Web, no 8-bit support — introduced locale
- 1999: World“s top-5 portal. We start with PostgreSQL 6.5.?
- Requirement: on the hardware ~ my smartphone to support > 1 mln. users/day
- Quickly run out of resources
- Denormalize, use arrays -> slow -> discover GiST → improve GiST - intarray with indexes
- Need FTS, made tsearch using intarray and GiST indexes
- Need fast search on hierarchical data — ltree — GiST indexes
- Need flexible schema — hstore — GiST index
- Need faster FTS — GIN index for tsearch, hstore
- Need misprint search — pg_trgm — GiST/GIN indexes
- Compete to NoSQL - better/binary json - jsonb — GIN index
- Need faster FTS — RUM access methods

2017: STILL USING PostgreSQL !

Postgres PROFESSIONAL

Extensibility makes PostgreSQL Universal Database !

PostgreSQL Forks (we love forks!)

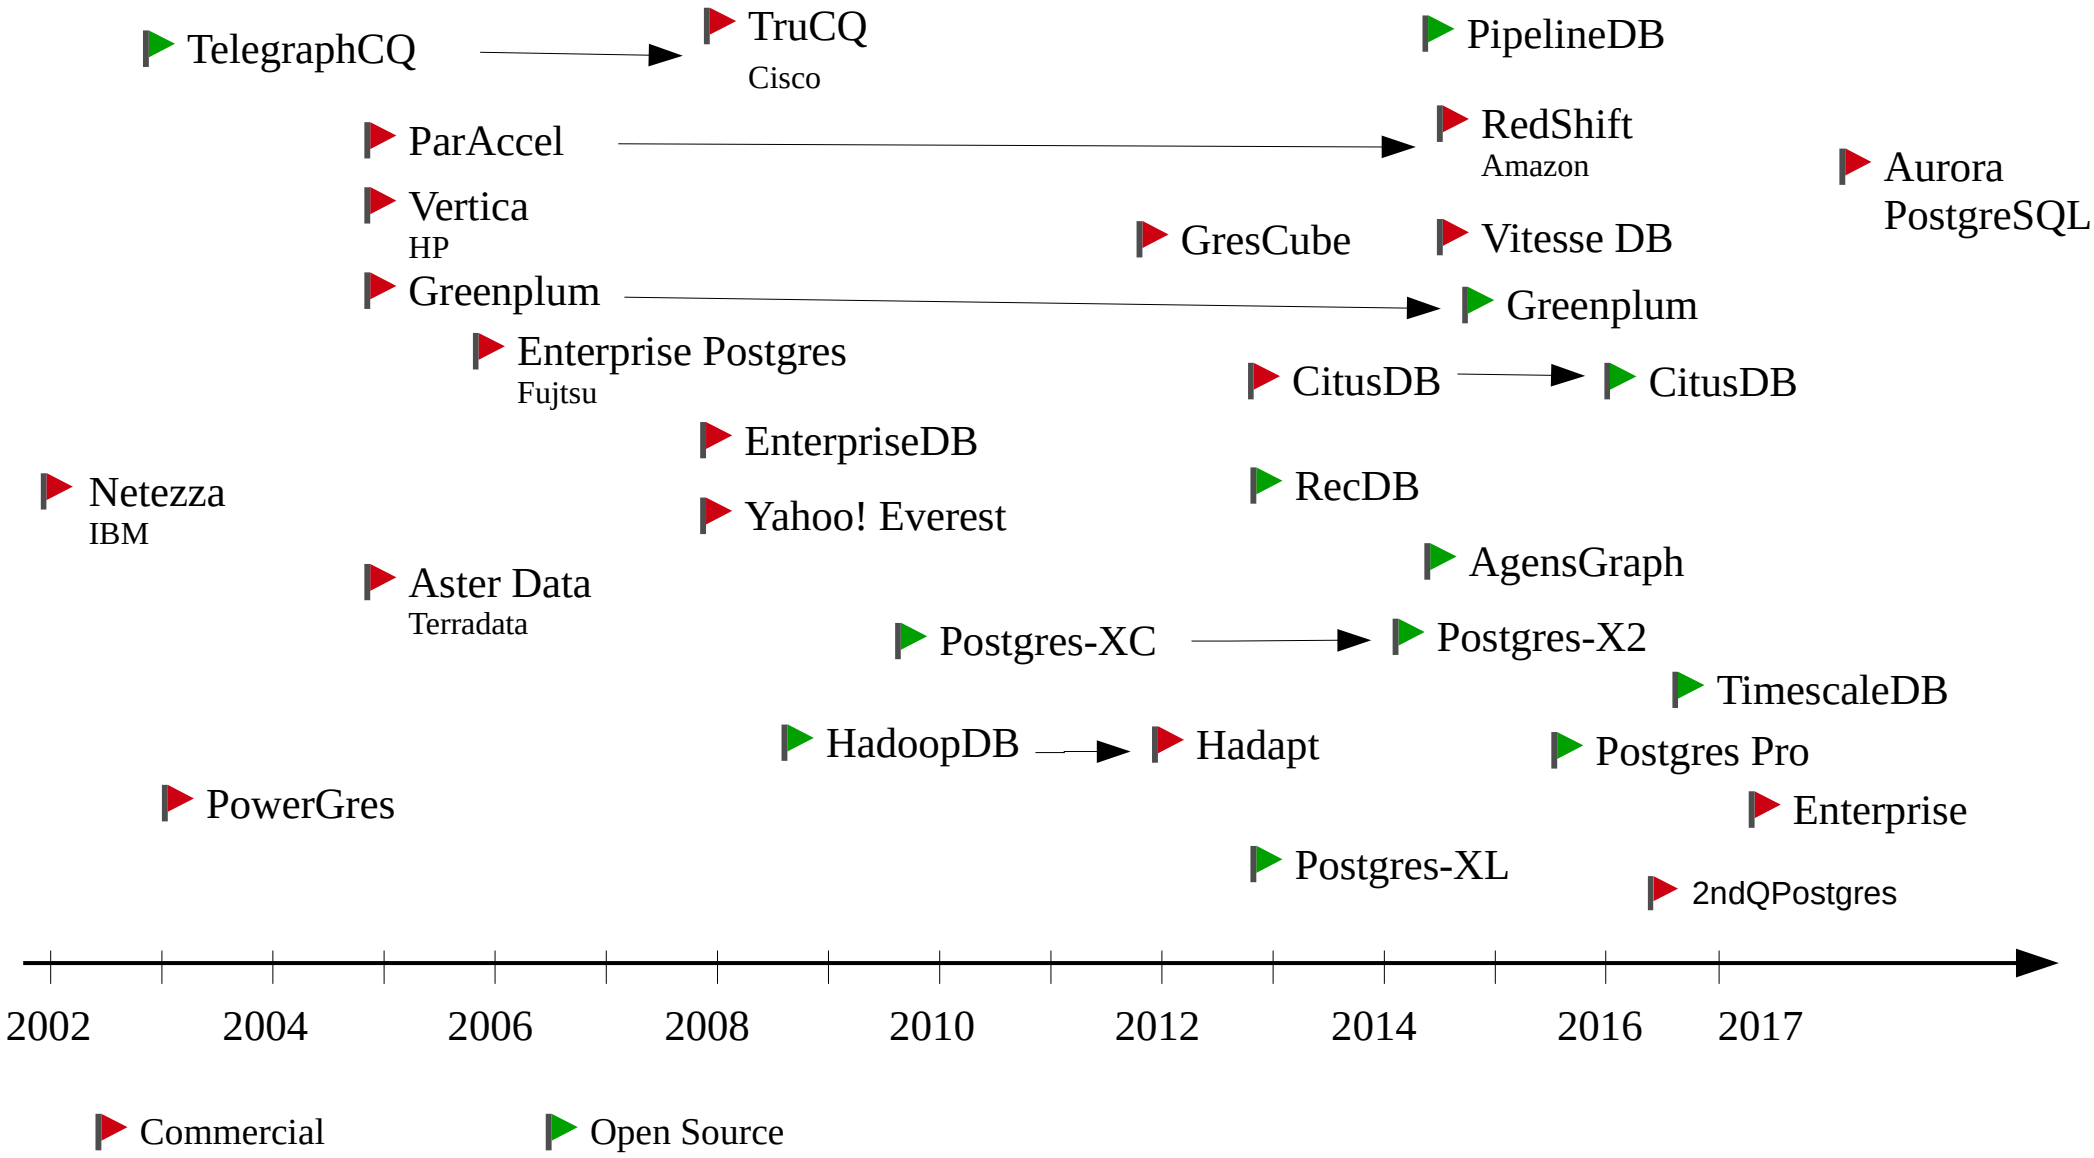

OLTP, MPP, OLAP, CLOUD, GIS, STREAM, NOSQL,TIMESERIES, BLOCKCHAIN, GPU

- PostgreSQL - OLTP, Web, Science
  extensions: GIS (PostGIS), time-series (TimeScaleDB), Stream data (PipelineDB), interactive analytics (Citus), vector ops, async mpp on GPU (PG-Strom)
- AgensGraph — Graph model, graph query language
- Postgres Pro[Enterpise] - Multimaster, Sharding, block compression, 64-bit TX, partitioning, adaptive planning, incremental backup, Credereum blockchain enabled DB
- Redshift — Cloud, column storage, compression, OLAP
- Greenplum - MPP, OLAP (complex reports)
- Postgres-XL — MPP OLAP (ACID)
- Vitesse DB — OLAP (JIT, column storage, threads)
- Aurora PostgreSQL — OLTP, S3 storage
- EDB Postgres Advanced Server - Oracle compatibility

Case 2: NoSQL Postgres

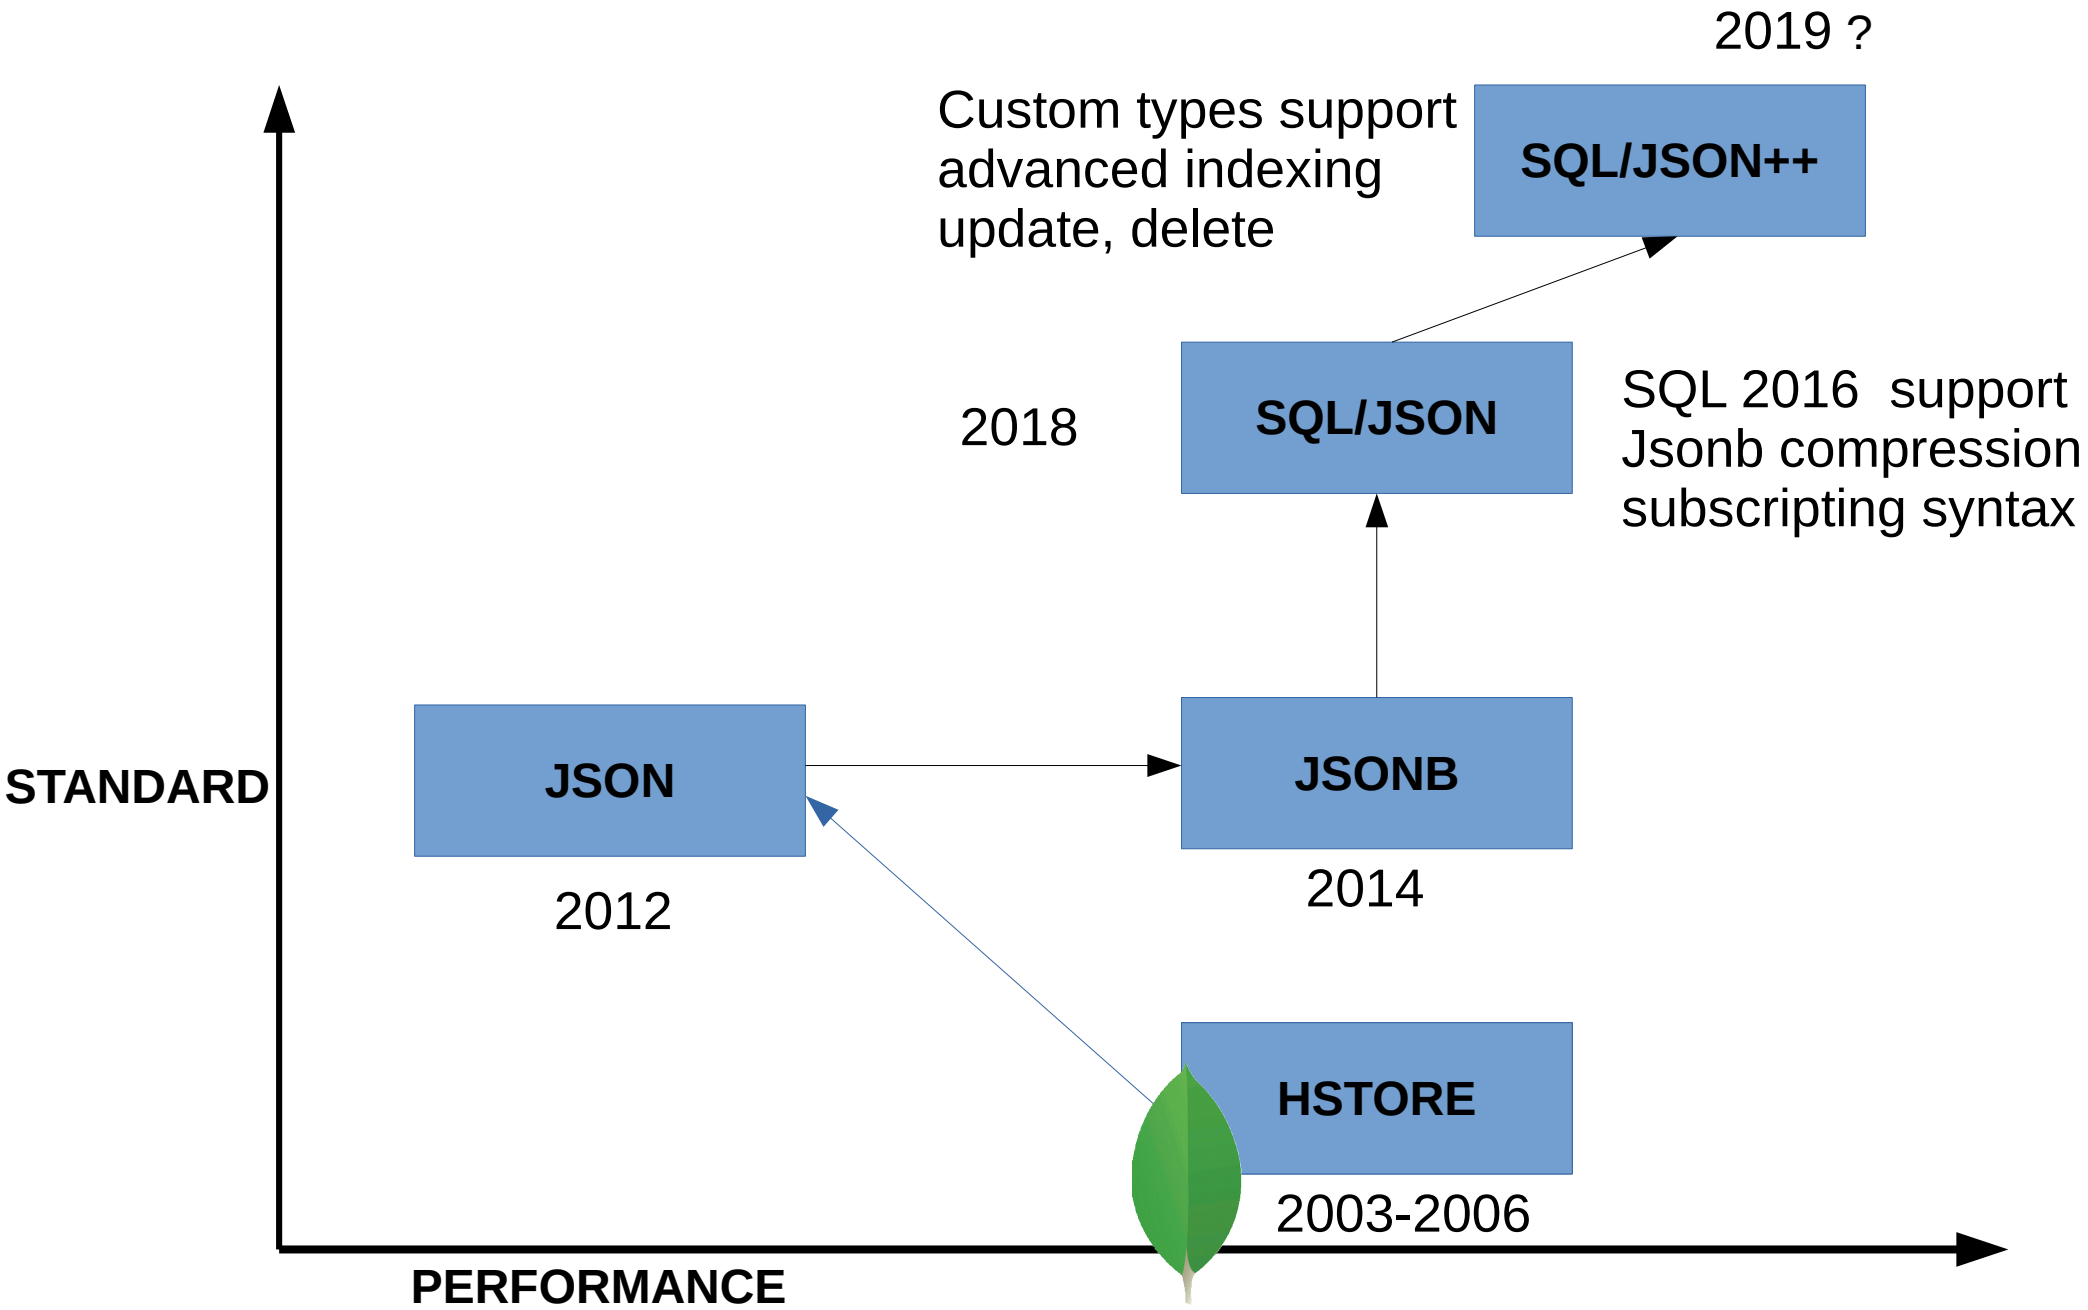

Postgres PROFESSIONAL

PostgreSQL users

+BIG RUSSIAN Enterprise !

Open source vs commercial DBMS

Popularity trend

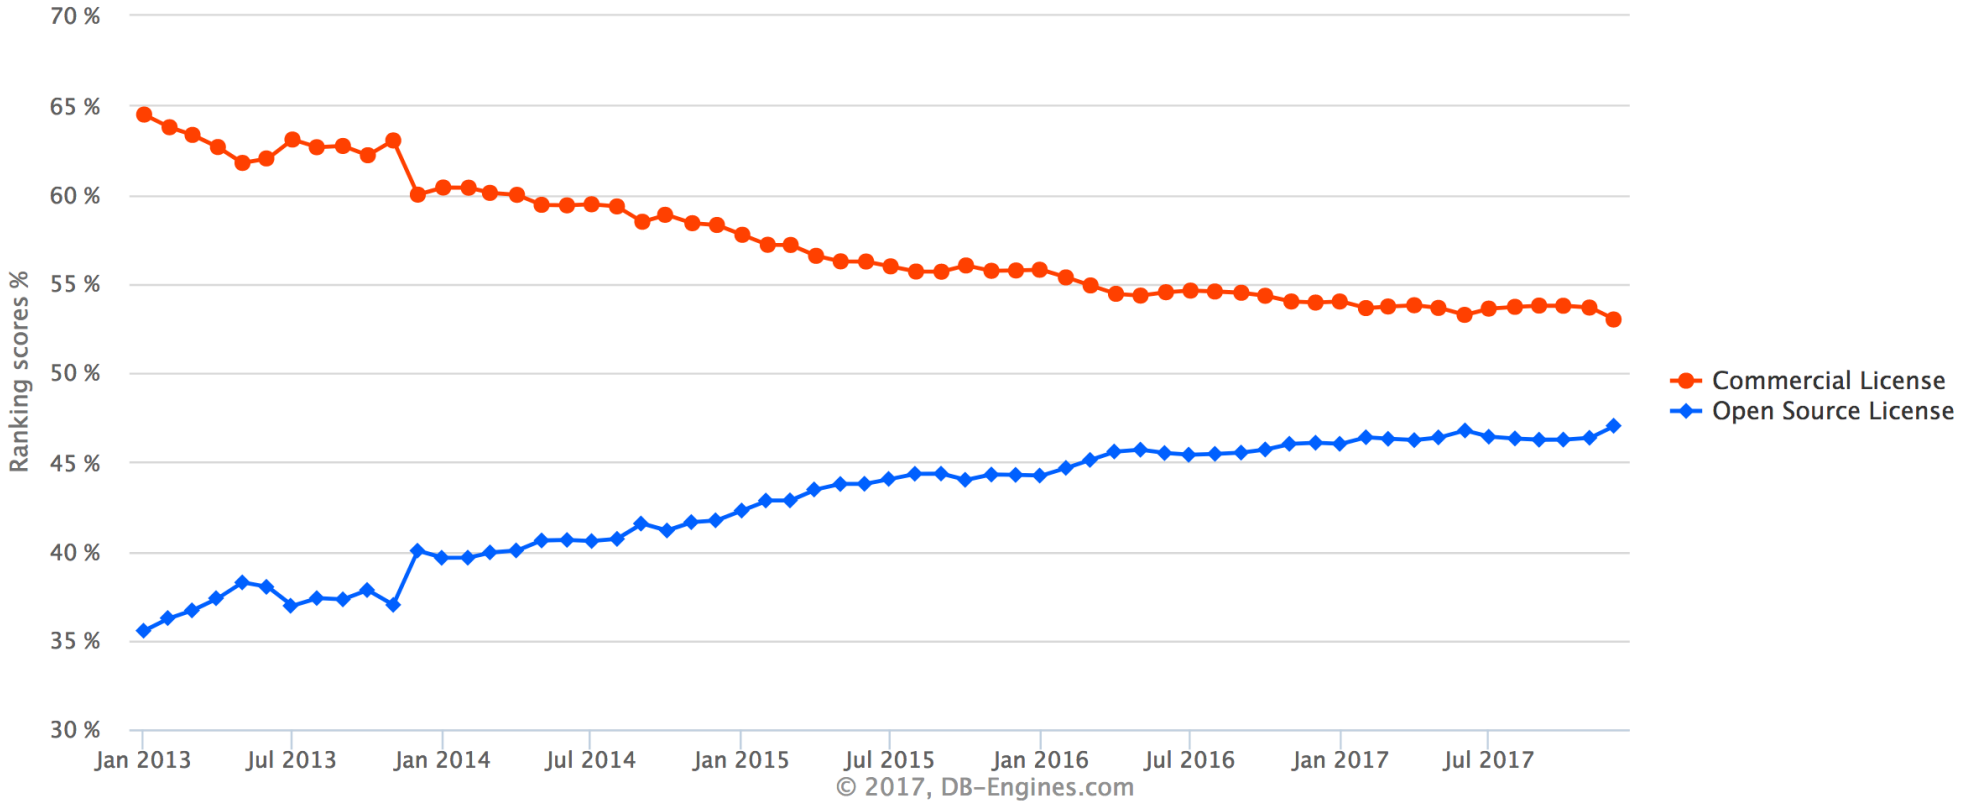

DBMS of the Year (2017): PostgreSQL

Previous winners of the DB-Engines DBMS of the Year Award:

PostgreSQL	2017
Microsoft SQL Server	2016
Oracle	2015
MongoDB	2014
MongoDB	2013

341 systems in ranking, January 2018

Rank			DBMS	Database Model	Score		
Jan 2018	Dec 2017	Jan 2017			Jan 2018	Dec 2017	Jan 2017
1.	1.	1.	Oracle	Relational DBMS	1341.94	+0.40	-74.78
2.	2.	2.	MySQL	Relational DBMS	1299.71	-18.36	-66.58
3.	3.	3.	Microsoft SQL Server	Relational DBMS	1148.07	-24.42	-72.89
4.	4.	↑ 5.	PostgreSQL	Relational DBMS	386.18	+0.75	+55.81
5.	5.	↓ 4.	MongoDB	Document store	330.95	+0.18	-0.96
6.	6.	6.	DB2	Relational DBMS	190.28	+0.70	+7.78
7.	7.	↑ 8.	Microsoft Access	Relational DBMS	126.70	+0.82	-0.75
8.	↑ 9.	↓ 7.	Cassandra	Wide column store	123.88	+0.67	-12.57
9.	↓ 8.	9.	Redis	Key-value store	123.14	-0.10	+4.44
10.	10.	↑ 11.	Elasticsearch	Search engine	122.55	+2.77	+16.38

https://db-engines.com/en/ranking

PROFESSIONAL Postgres

Hacker News Hiring Trends

October 2017 Hacker News Hiring Trends

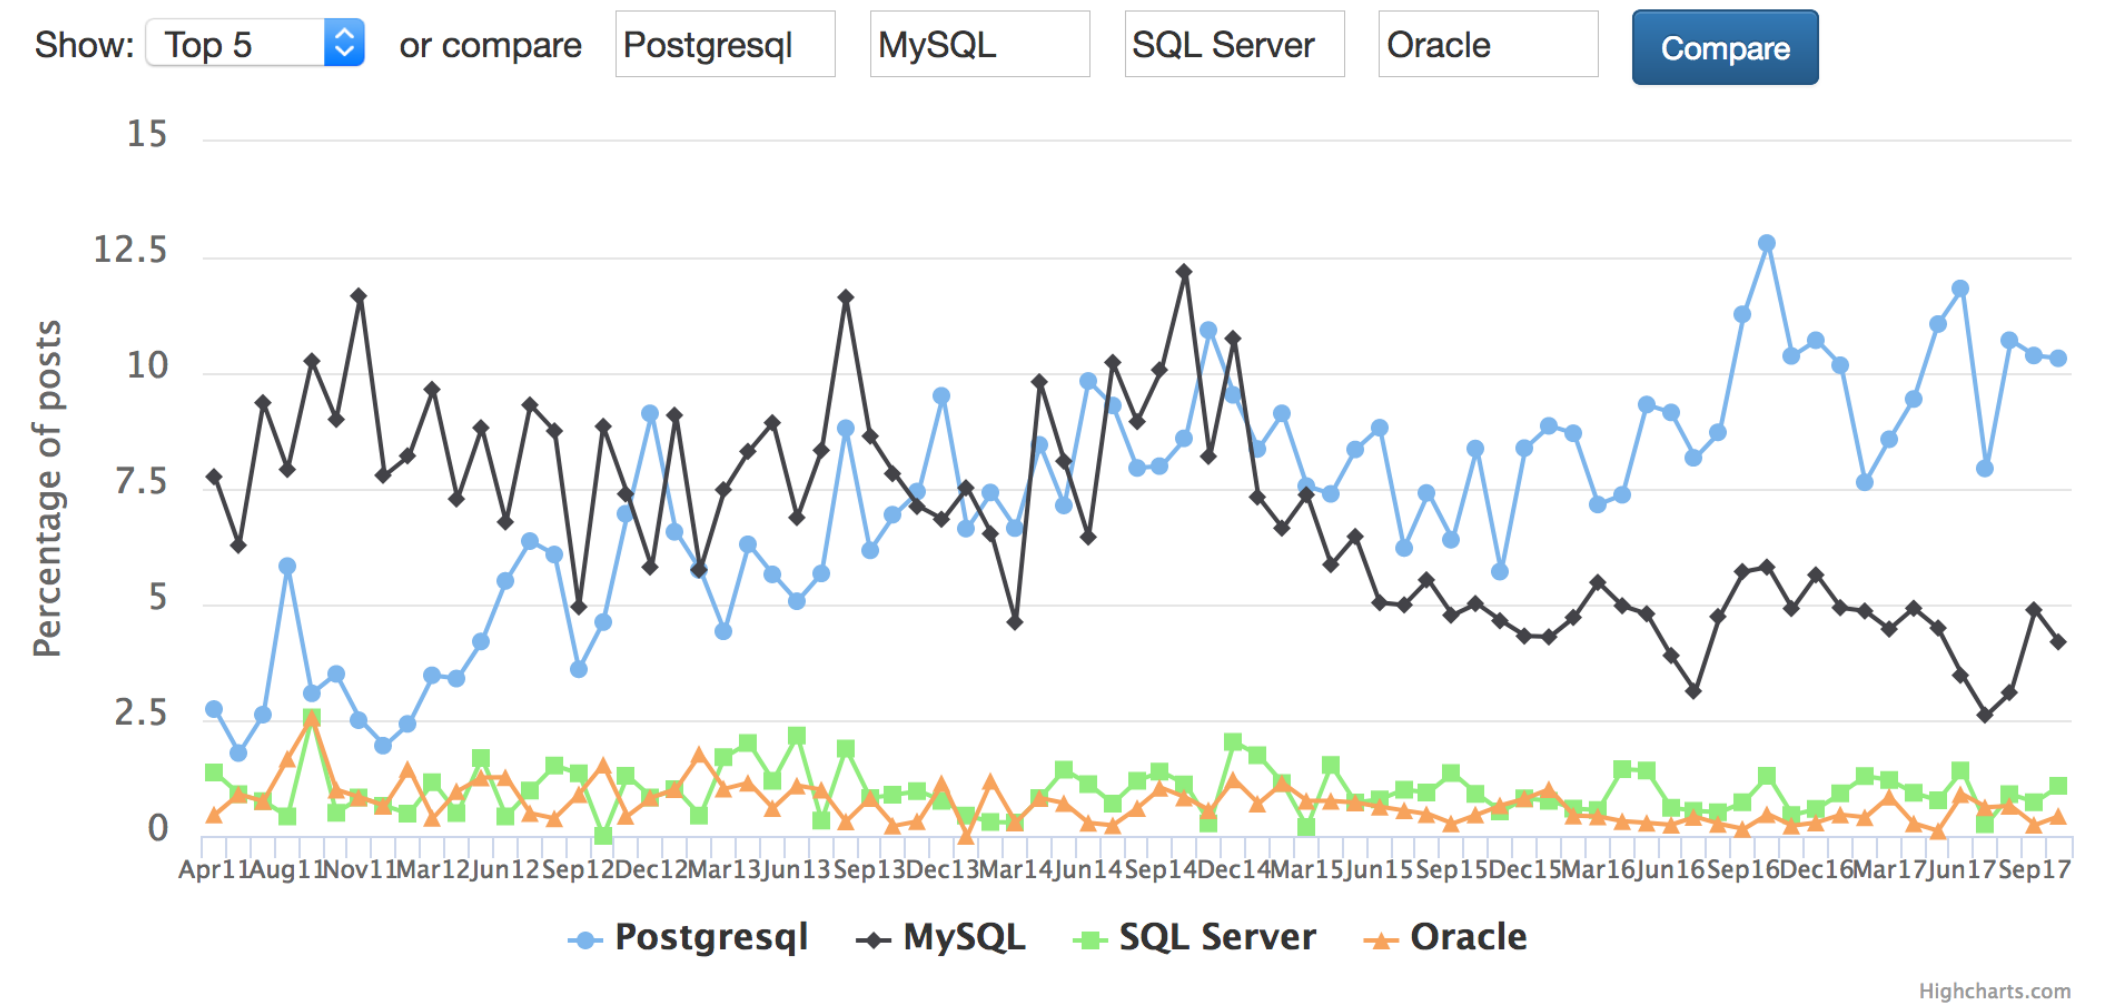

PostgreSQL in Russia is #1

PostgreSQL reach Oracle in Russia

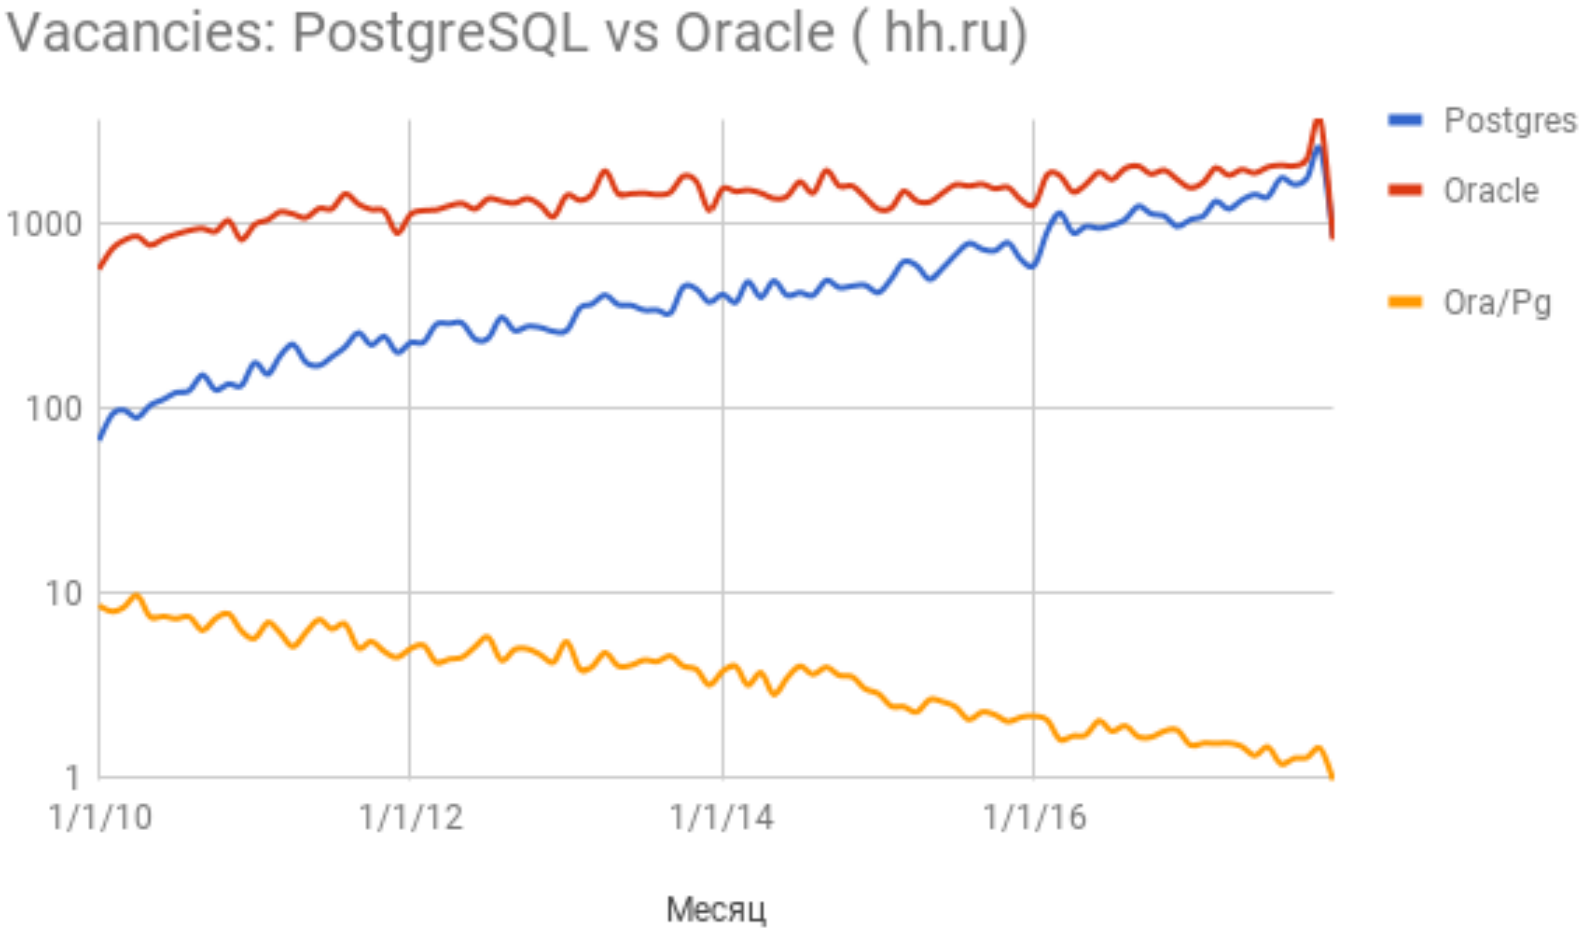

Case 3: PostgreSQL and Supernovae

SN 1987a, Type II, +2.9, Tarantula Nebulae in LMC,
168 000 ly, progenitor: Sanduleak -69° 202,

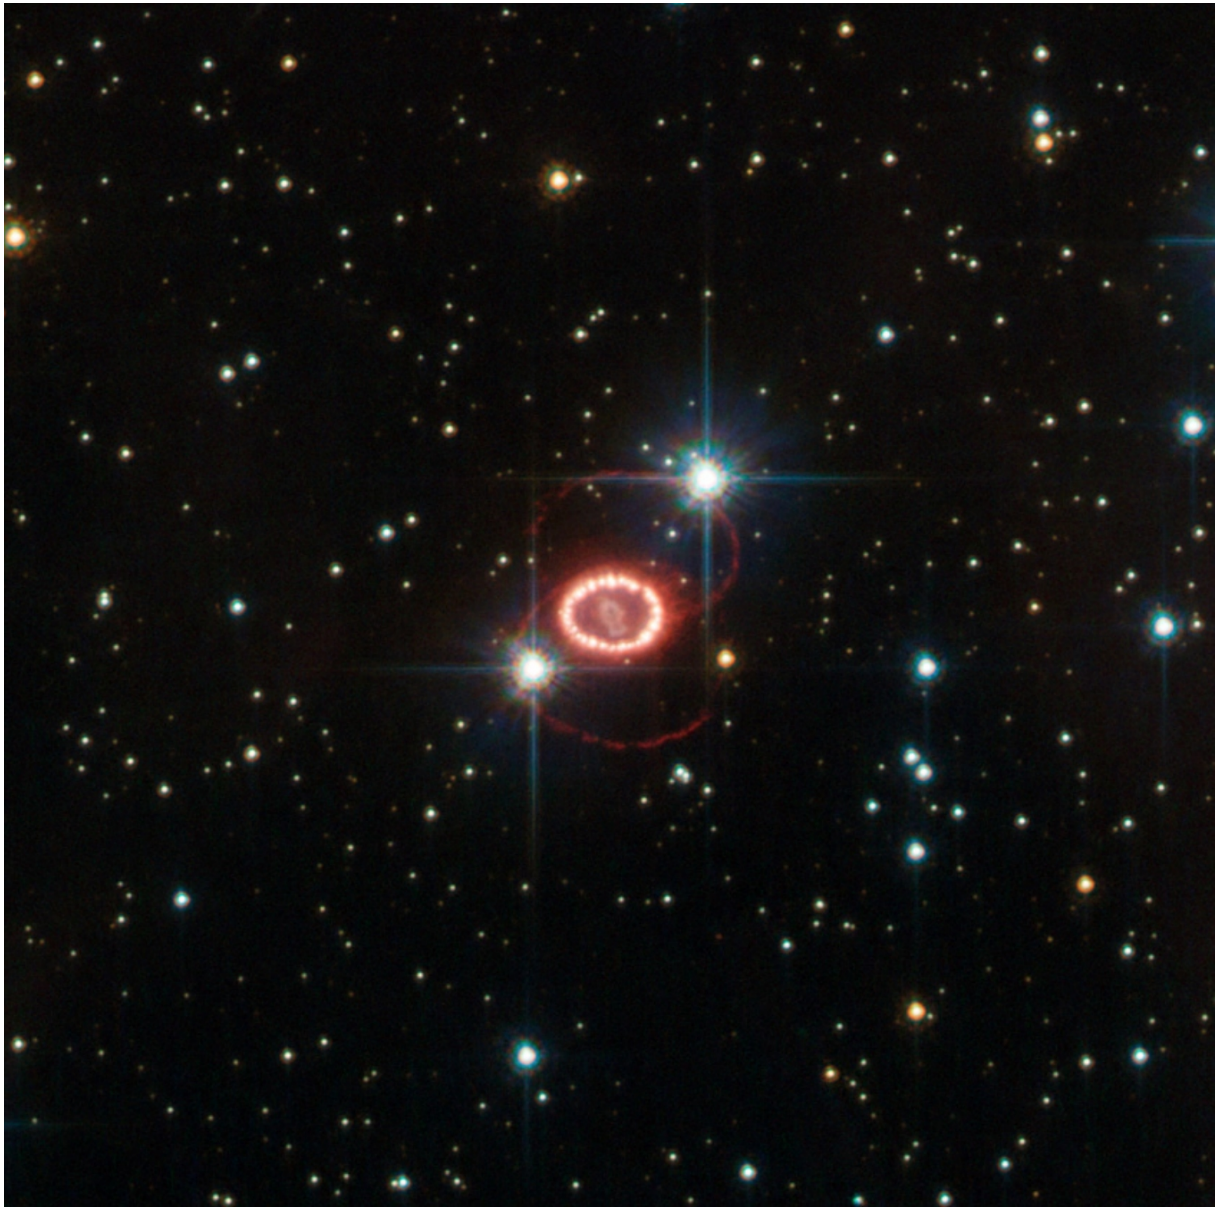

The Scale of the Universe

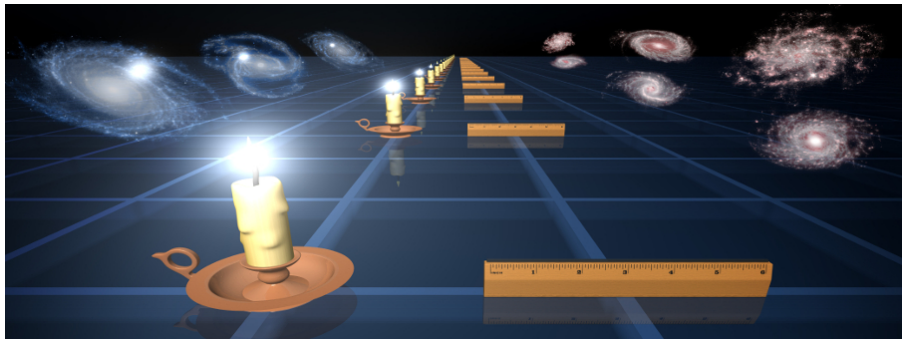

- Supernovae(Ia) - «standard candles»
- Used to measure the distance to the host galaxy

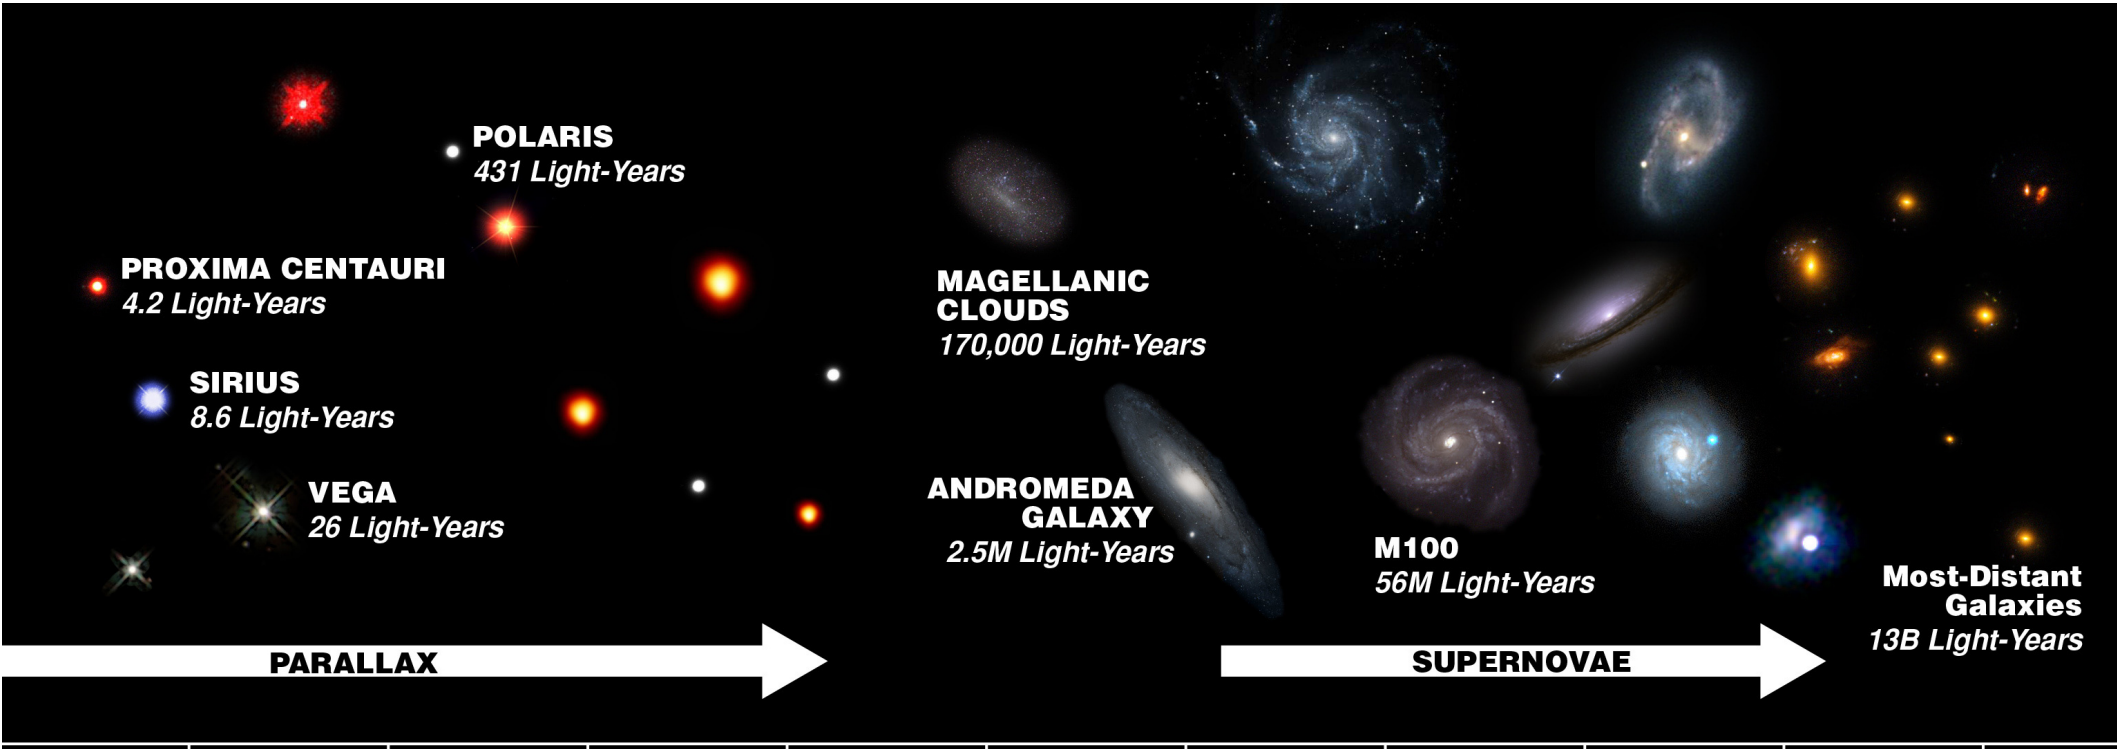

10	100	1 000	10 000	100 000	1 000 000	10 000 000	100 000 000	1 000 000 000	10 000 000 000	BIG BANG

Distance in Light-Years

CEPHEIDS

REDSHIFT

What is the fate of the Universe ?

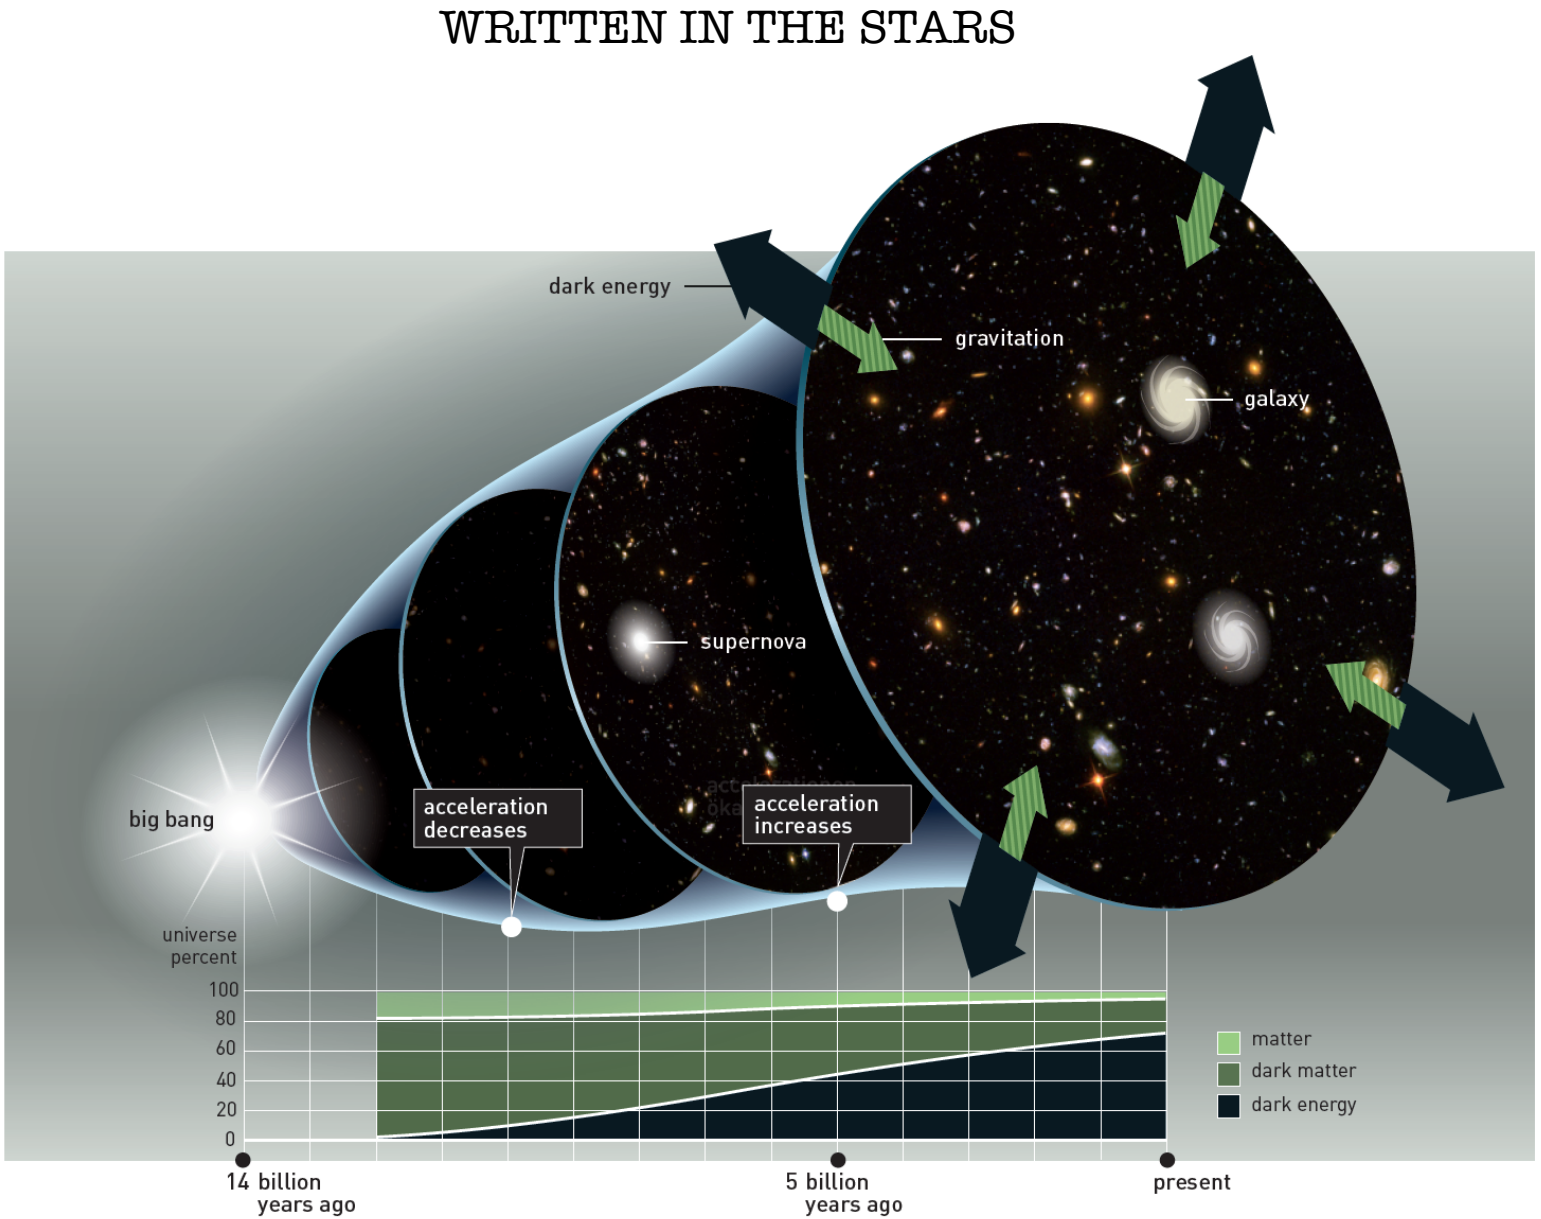

Figure 1. The world is growing. The expansion of the Universe began with the Big Bang 14 billion years ago, but slowed down during the first several billion years. Eventually it started to accelerate. The acceleration is believed to be driven by dark energy, which in the beginning constituted only a small part of the Universe. But as matter got diluted by the expansion, the dark energy became more dominant.

Change photo

Oleg Bartunov

Edit | Follow

Lomonosov Moscow State University
astrophysics, databases, astroinformatics
Verified email at sai.msu.su - Homepage
My profile is public

Google Scholar

Citation indices	All	Since 2012
Citations	1408	362
h-index	13	9
i10-index	17	9

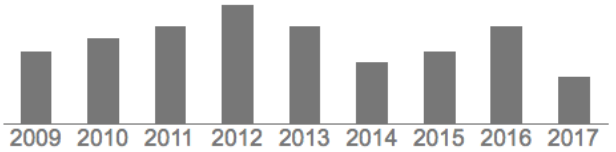

Add co-authors

Sergei B. Popov

Co-authors Edit...

Massimo Turatto
Mario Hamuy
Sergey Koposov
Stefano Benetti
Sergey Karpov
Igor Chilingarian

Title	Add	More	1–20	Cited by	Year
The rate of supernovae from the combined sample of five searches E Cappellaro, M Turatto, DY Tsvetkov, OS Bartunov, C Pollas, R Evans, ... arXiv preprint astro-ph/9611191				406	1996
Radiation hydrodynamics of SN 1987A. I. Global analysis of the light curve for the first 4 months S Blinnikov, P Lundqvist, O Bartunov, K Nomoto, K Iwamoto The Astrophysical Journal 532 (2), 1132				179	2000
A comparative modeling of supernova 1993J SI Blinnikov, R Eastman, OS Bartunov, VA Popolitov, SE Woosley The Astrophysical Journal 496 (1), 454				159	1998
The rate of supernovae. II. The selection effects and the frequencies per unit blue luminosity E Cappellaro, M Turatto, DY Tsvetkov, OS Bartunov, IN Makarova arXiv preprint astro-ph/9302017				118	1993
Non-equilibrium radiative transfer in supernova theory-models of linear Type-II supernovae SI Blinnikov, OS Bartunov Astronomy and Astrophysics 273, 106				102	1993
Distribution of supernovae relative to spiral arms and H II regions OS Bartunov, DY Tsvetkov, IV Filimonova Publications of the Astronomical Society of the Pacific 106 (706), 1276				100	1994

M31 (Andromeda), AZT-5 telescope

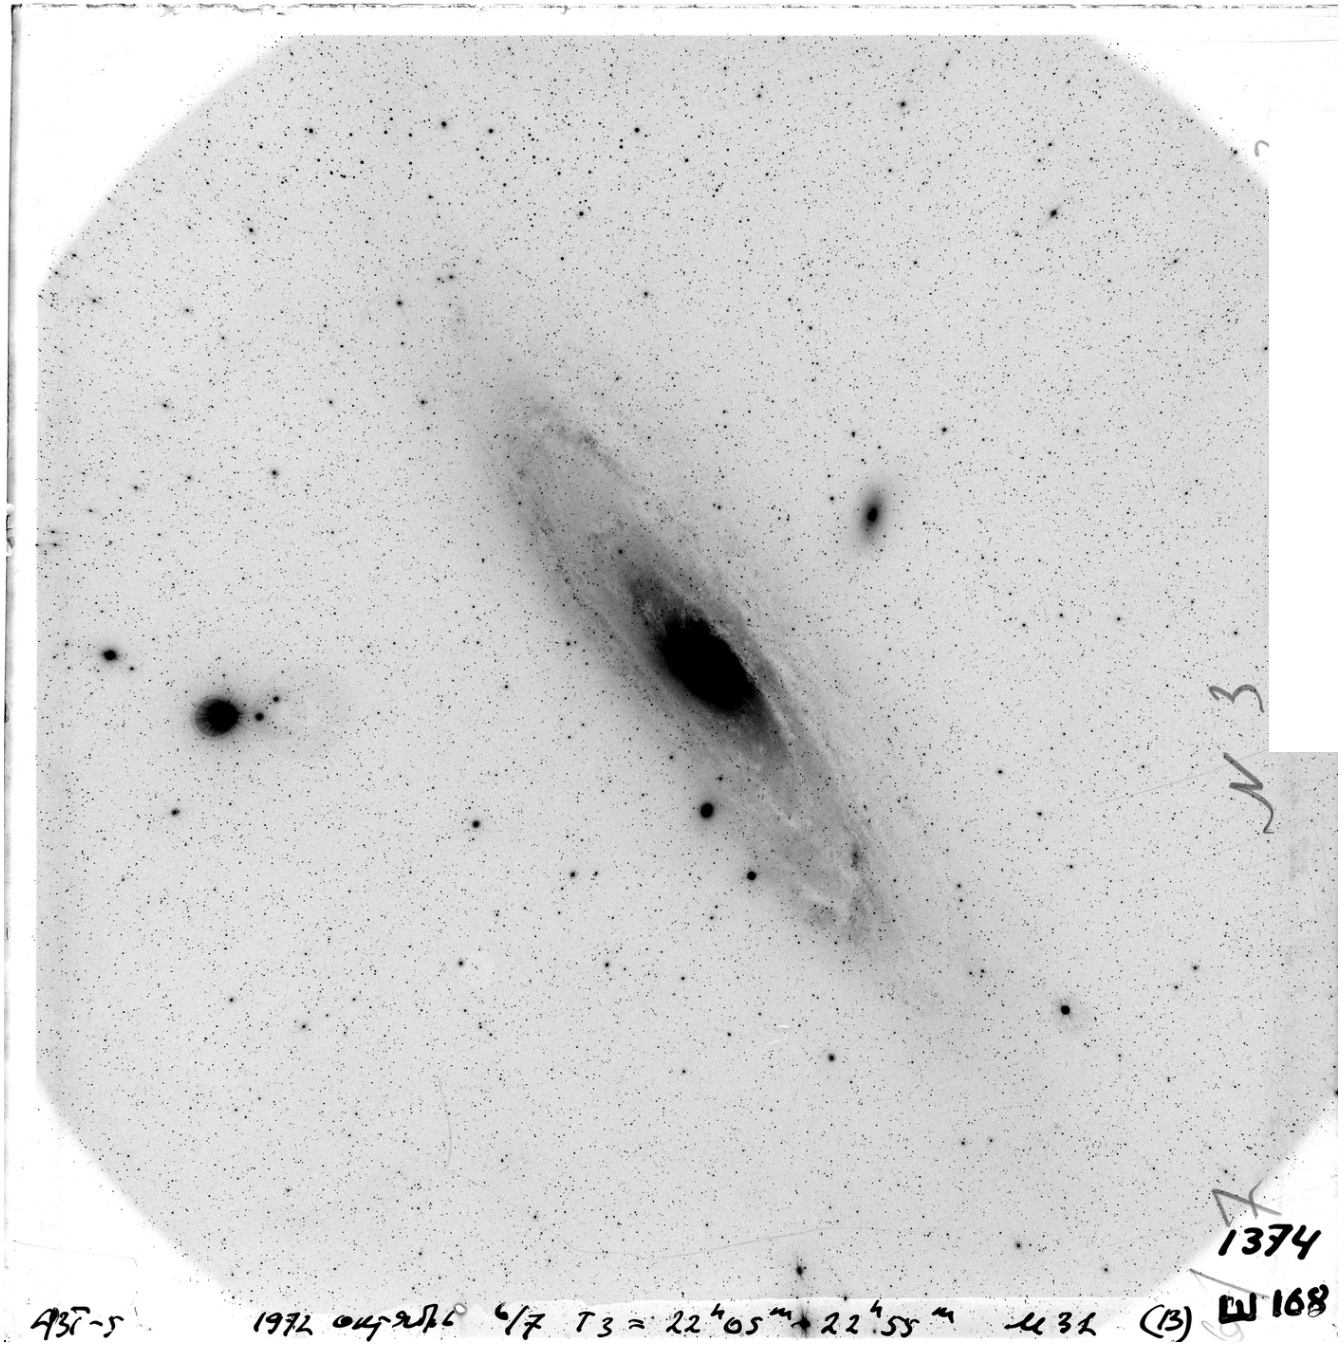

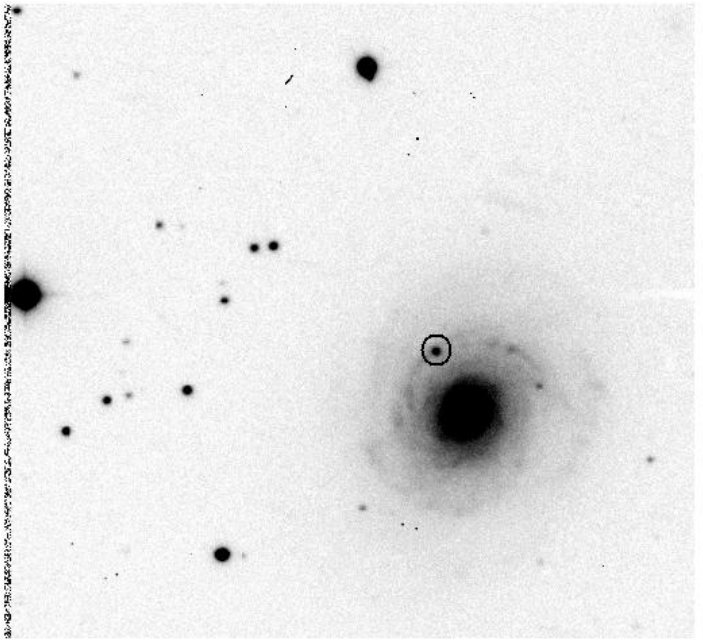

SN 2008fv in NGC 3147, Draco
Dmitry Tsvetkov, SAI MSU

Blink Comparator (manual discovery)
Many hours of hard work !

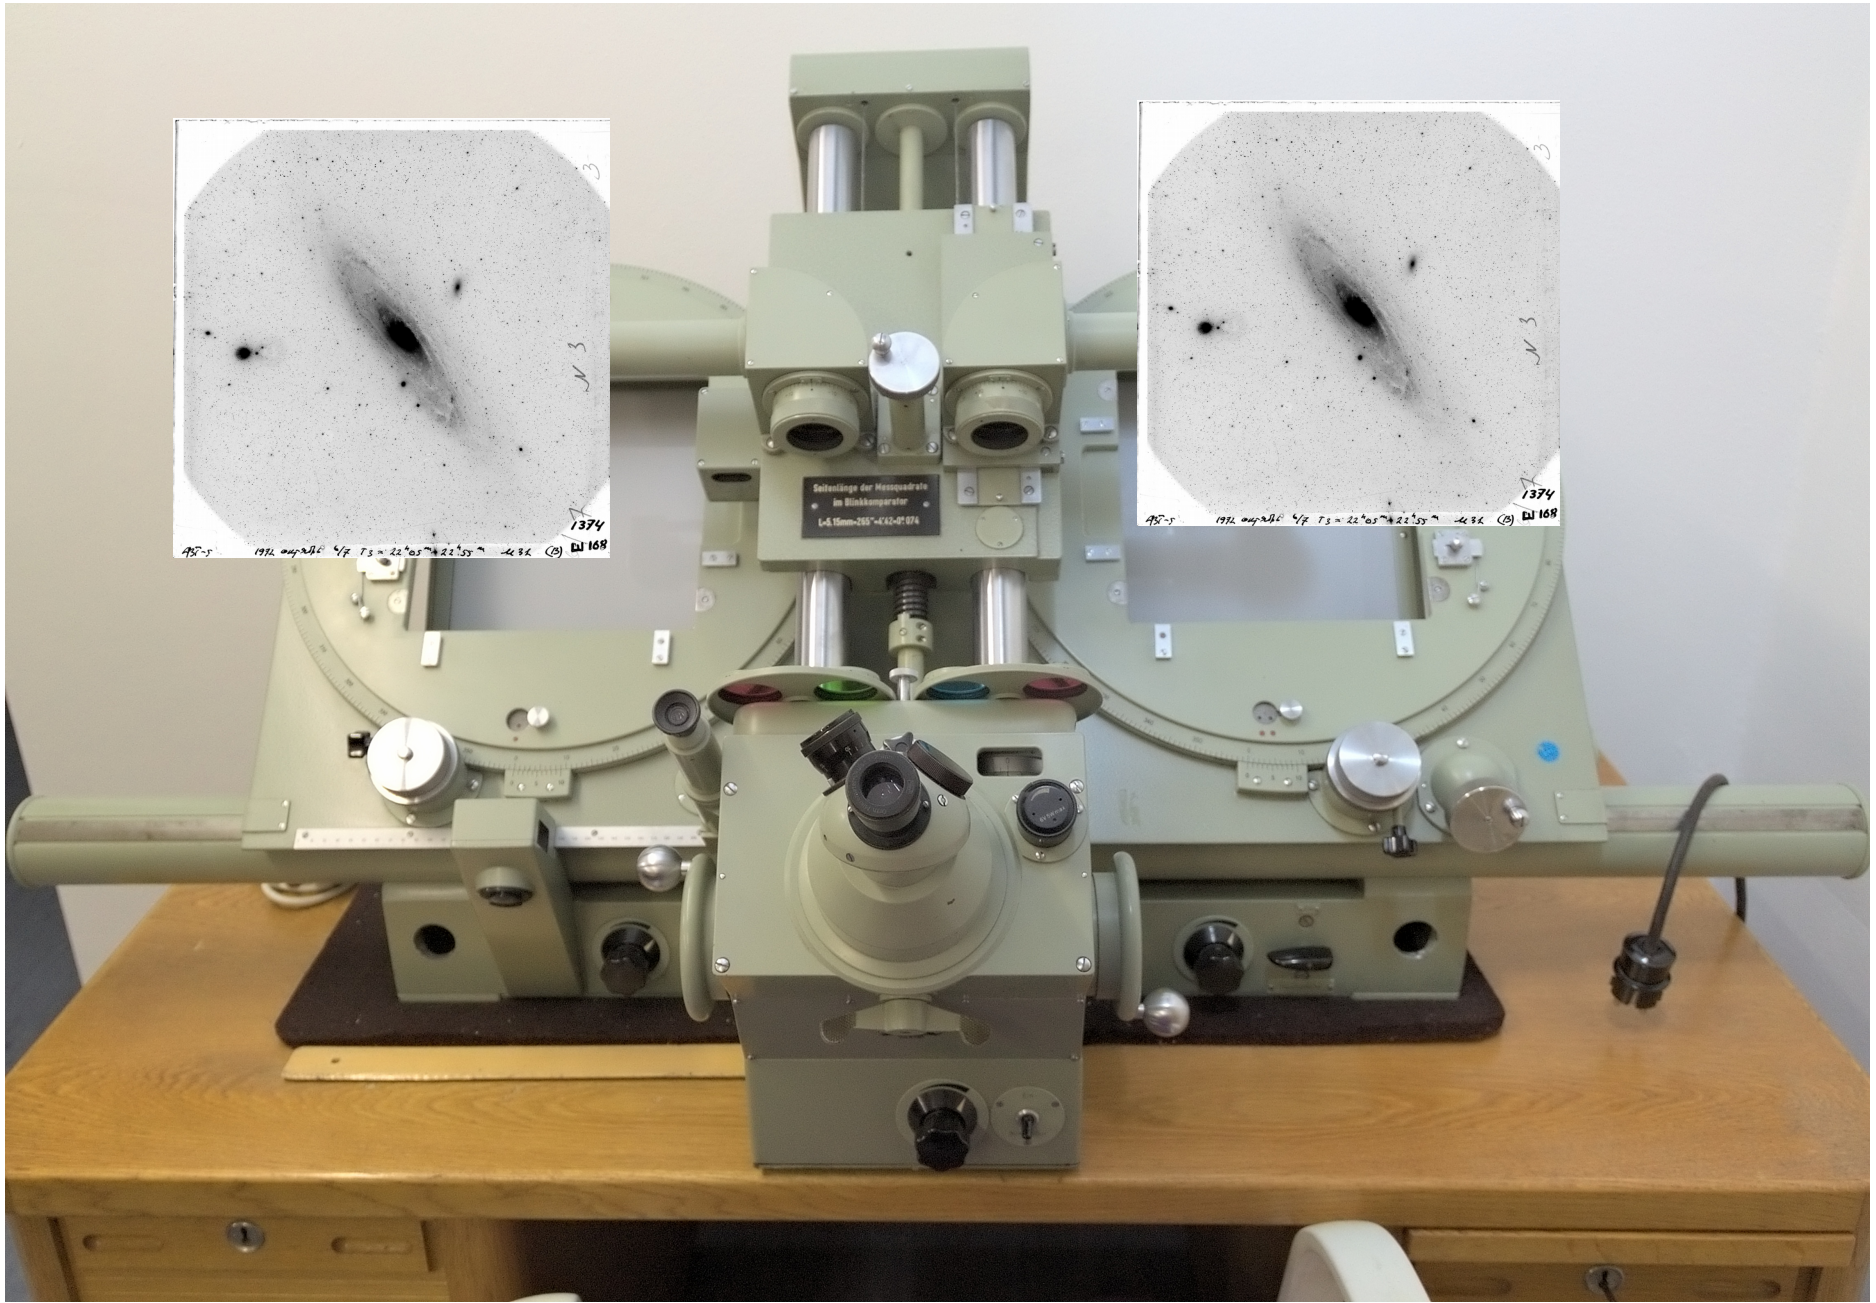

M31 (Andromeda), AZT-5 telescope

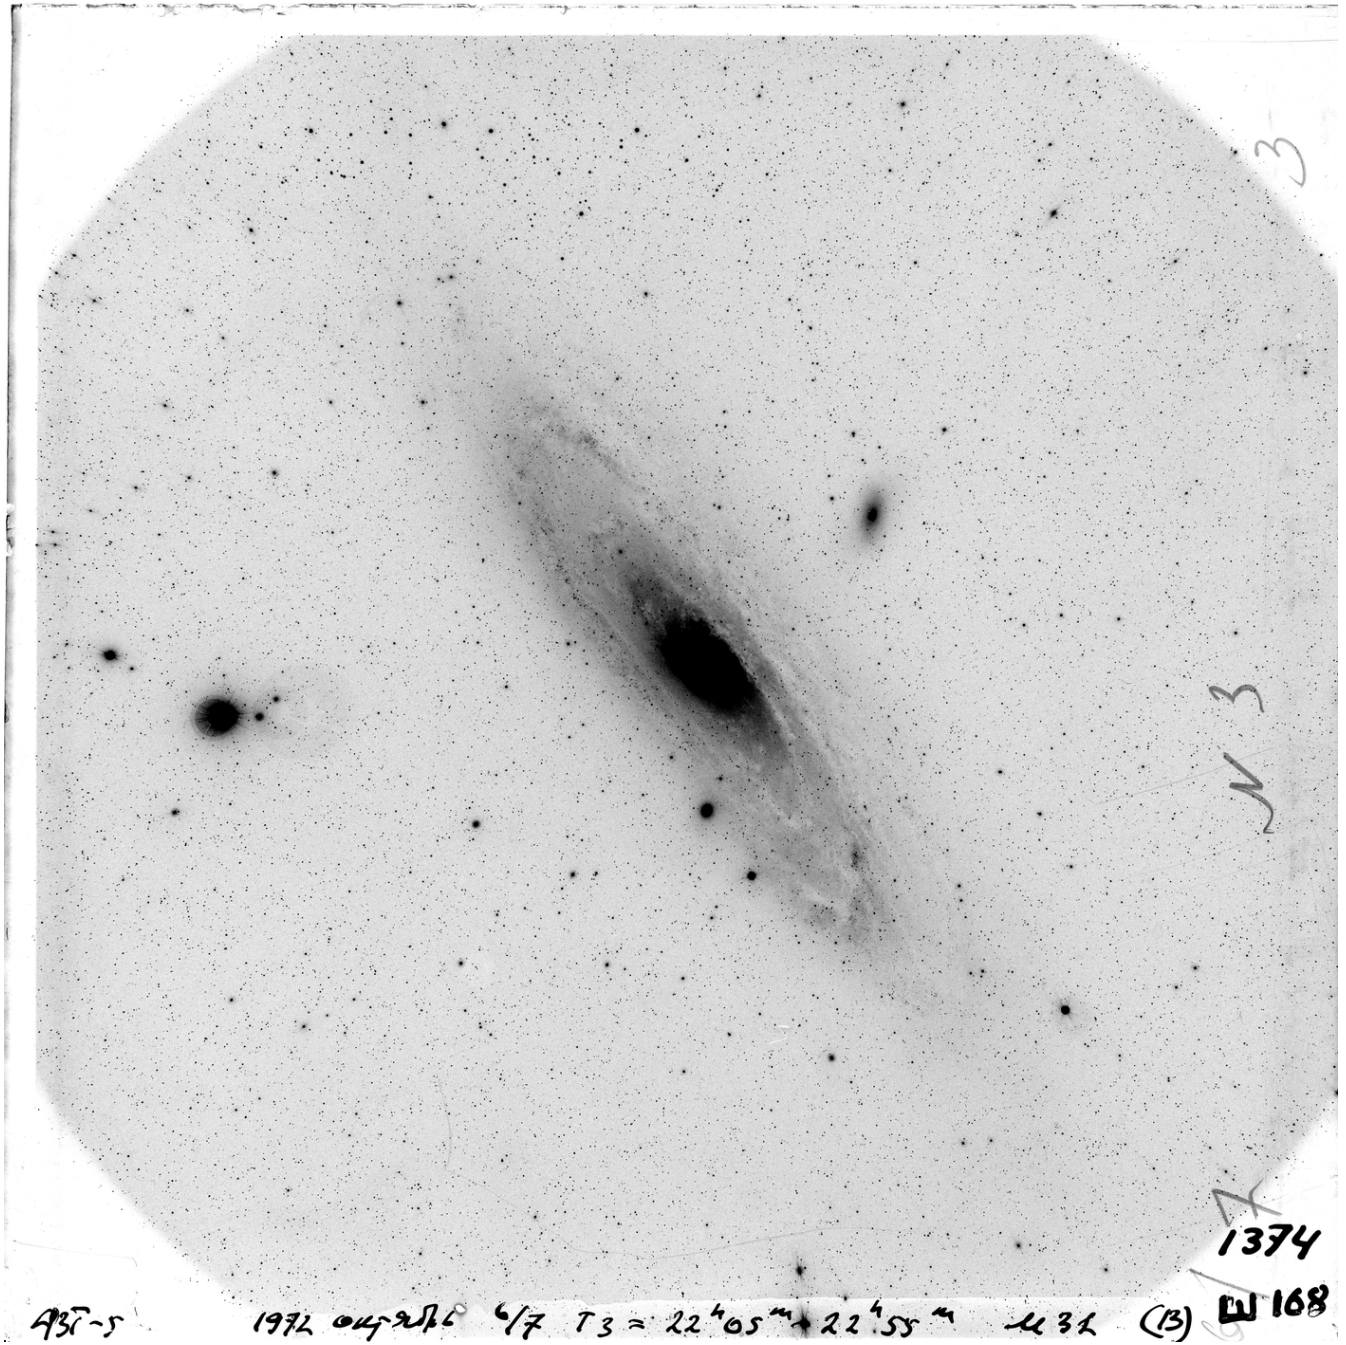


```
x-00001|1.1458447|-89.9186147
x-00002|1.3300139|-89.9332336
x-00003|3.2556022|-89.9641031
x-00004|3.6464625|-89.9060142
x-00005|6.3110253|-89.9523947
x-00006|6.6275517|-89.9279197
x-00007|7.8266025|-89.9129272
x-00008|9.0694378|-89.9714031
x-00009|9.6627953|-89.9244314
x-00010|10.0494292|-89.9705058
x-00011|10.4863922|-89.9699058
x-00012|11.0953692|-89.9016031
x-00013|11.3240233|-89.9344336
x-00014|11.7906064|-89.9070308
x-00015|12.0416581|-89.9300586
x-00016|12.0522308|-89.9002281
x-00017|12.2808536|-89.9107669
x-00018|13.0316142|-89.9214558
x-00019|13.8727033|-89.9577031
x-00020|14.6546639|-89.9191919
x-00021|18.3035981|-89.9447475
x-00022|18.5185631|-89.9446836
x-00023|19.8675597|-89.9836308
x-00024|20.9699533|-89.9226864
x-00025|21.6777744|-89.9256808
x-00026|23.3660669|-89.9036558
x-00027|24.2841308|-89.9516475
x-00028|24.3273161|-89.9202392
x-00029|24.5540458|-89.9246003
x-00030|24.5655172|-89.9122336
x-00031|26.3487519|-89.9460336
x-00032|26.5268008|-89.9311503
x-00033|26.6070808|-89.9271808
x-00034|27.4104919|-89.9768558
x-00035|27.8290442|-89.9304622
x-00036|28.5552036|-89.9199117
x-00037|29.4407347|-89.9762836
x-00038|30.5729608|-89.9377753
x-00039|30.7101131|-89.9105642
x-00040|33.2918250|-89.9106614
x-00041|33.4843678|-89.9442058
```

Spatial Join (Machine Discovery, < 1s)

Observations: 10^5

```
x-00001|1.1458447|-89.9186147
x-00002|1.3300139|-89.9332336
x-00003|3.2556022|-89.9641031
x-00004|3.6464625|-89.9060142
x-00005|6.3110253|-89.9523947
x-00006|6.6275517|-89.9279197
x-00007|7.8266025|-89.9129272
x-00008|9.0694378|-89.9714031
x-00009|9.6627953|-89.9244314
x-00010|10.0494292|-89.9705058
x-00011|10.4863922|-89.9699058
x-00012|11.0953692|-89.9016031
x-00013|11.3240233|-89.9344336
x-00014|11.7906064|-89.9070308
x-00015|12.0416581|-89.9300586
x-00016|12.0522308|-89.9002281
x-00017|12.2808536|-89.9107669
x-00018|13.0316142|-89.9214558
x-00019|13.8727033|-89.9577031
x-00020|14.6546639|-89.9191919
x-00021|18.3035981|-89.9447475
x-00022|18.5185631|-89.9446836
x-00023|19.8675597|-89.9836308
x-00024|20.9699533|-89.9226864
x-00025|21.6777744|-89.9256808
x-00026|23.3660669|-89.9036558
x-00027|24.2841308|-89.9516475
x-00028|24.3273161|-89.9202392
x-00029|24.5540458|-89.9246003
x-00030|24.5655172|-89.9122336
x-00031|26.3487519|-89.9460336
x-00032|26.5268008|-89.9311503
x-00033|26.6070808|-89.9271808
x-00034|27.4104919|-89.9768558
x-00035|27.8290442|-89.9304622
x-00036|28.5552036|-89.9199117
x-00037|29.4407347|-89.9762836
x-00038|30.5729608|-89.9377753
```

$\bowtie\theta$

Catalog(s): 10^9

```
t-0000001|1.1458447|-89.9186147|0.015|0.028
t-0000002|1.3300139|-89.9332336|0.050|0.110
t-0000003|3.2556022|-89.9641031|0.050|0.050
t-0000004|3.6464625|-89.9060142|0.204|0.224
t-0000005|6.3110253|-89.9523947|0.114|0.050
t-0000006|6.6275517|-89.9279197|0.098|0.150
t-0000007|7.8266025|-89.9129272|0.025|0.021
t-0000008|9.0694378|-89.9714031|0.200|0.200
t-0000009|9.6627953|-89.9244314|0.000|0.000
t-0000010|10.0494292|-89.9705058|0.050|0.228
t-0000011|10.4863922|-89.9699058|0.200|0.200
t-0000012|11.0953692|-89.9016031|0.050|0.259
t-0000013|11.3240233|-89.9344336|0.050|0.050
t-0000014|11.7906064|-89.9070308|0.159|0.131
t-0000015|12.0416581|-89.9300586|0.216|0.050
t-0000016|12.0522308|-89.9002281|0.050|0.050
t-0000017|12.2808536|-89.9107669|0.050|0.050
t-0000018|13.0316142|-89.9214558|0.152|0.120
t-0000019|13.8727033|-89.9577031|0.050|0.121
t-0000020|14.6546639|-89.9191919|0.050|0.069
t-0000021|18.3035981|-89.9447475|0.139|0.440
t-0000022|18.5185631|-89.9446836|0.057|0.268
t-0000023|19.8675597|-89.9836308|0.050|0.120
t-0000024|20.9699533|-89.9226864|0.050|0.050
t-0000025|21.6777744|-89.9256808|0.055|0.105
t-0000026|23.3660669|-89.9036558|0.050|0.135
t-0000027|24.2841308|-89.9516475|0.213|0.050
t-0000028|24.3273161|-89.9202392|0.550|0.999
t-0000029|24.5540458|-89.9246003|0.160|0.086
t-0000030|24.5655172|-89.9122336|0.205|0.050
t-0000031|26.3487519|-89.9460336|0.050|0.095
t-0000032|26.5268008|-89.9311503|0.335|0.245
t-0000033|26.6070808|-89.9271808|0.050|0.075
t-0000034|27.4104919|-89.9768558|0.094|0.090
t-0000035|27.8290442|-89.9304622|0.017|0.019
t-0000036|28.5552036|-89.9199117|0.050|0.115
t-0000037|29.4407347|-89.9762836|0.635|0.265
t-0000038|30.5729608|-89.9377753|0.314|0.170
```

Astronomy meets database

Indexing the SKY with PostgreSQL

HEALPIX

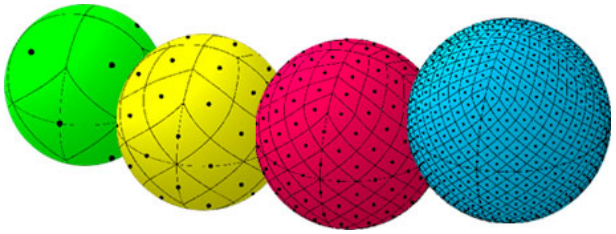

HTM

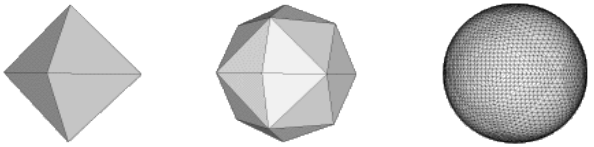

Q3C (PG)

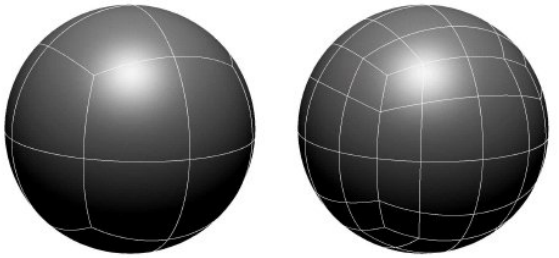

The sphere segmentation in Q3C

PGSphere (PG)

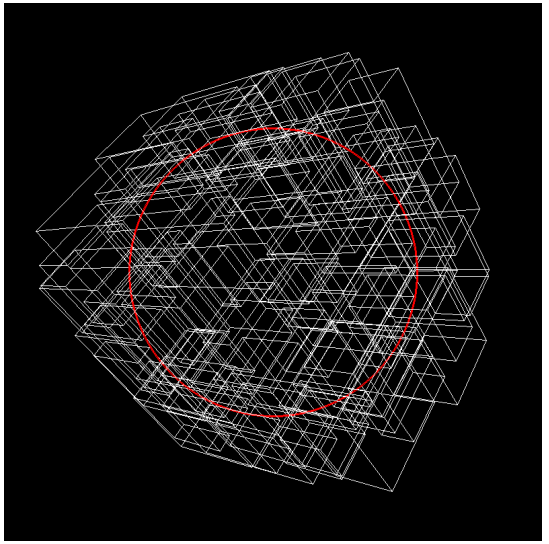

WSDB Whole Sky DataBase

Database of astronomic catalogues in Cambridge University

~5 dbs,~ 40 users, up to ~ 10^7 queries per day, size 40Tb

pg 9.4 + q3c + hstore

Example of research:Koposov, S. E., Belokurov, V., Torrealba, G., & Evans, N. W. (2015). Beasts of the Southern Wild: Discovery of nine Ultra Faint satellites in the vicinity of the Magellanic Clouds. The Astrophysical Journal, 805(2), 130.

Gaia Alerts Database

Real time Detection of alerts in the Gaia

~10 dbs, 10 users, up to ~ 10^6 queries per day, size 30Tb

pg 9.3 + synchronous replication + q3c

Example of research:Campbell, H. C., Marsh, T. R., Fraser, M., Hodgkin, S. T., de Miguel, E., Gänsicke, B. T., ... & Koposov, S. E. (2015). Total eclipse of the heart: the AM CVn Gaia14aae/ASSASN-14cn. Monthly Notices of the Royal Astronomical Society, 452(1), 1060-1067.

MASTER database

Robotic network telescopes by SAI MSU

8 observatories (5 in Russia, 3 outside)

total size ~100TB

pg 9.0-9.4 + pgsphere + replication

See:

- Lipunov, Vladimir, et al. "Master robotic net." Advances in Astronomy 2010 (2010).
- Kornilov, Victor G., et al. "Robotic optical telescopes global network MASTER II. Equipment, structure, algorithms." Experimental Astronomy 33.1 (2012): 173-196.

PostgreSQL Future

Several Postgres groups are working on

Postgres Distributed

Postgres Parallel

Postgres Asynchronous

Postgres Extendable+

Postgres NoSQL

Postgres Scalable

Postgres Trusted

References

1. https://www.postgresql.org/ - community site
2. https://wiki.postgresql.org/ - wiki
   - https://wiki.postgresql.org/wiki/FAQ - PostgreSQL FAQ
3. https://www.postgresql.org/list/ - mailing lists archive
   - https://postgrespro.com/list/ - search mailing lists
4. https://planet.postgresql.org/ - Planet PostgreSQL (blogs)
5. https://www.postgresql.org/about/events/ - Events
   PGConf.eu, PGConf.ru, PGConf.asia, PGCon.org, PGConf.org — major conferences
6. https://postgrespro.ru/education - Postgres Professional education materials
7. https://pgxn.org/ - PostgreSQL Extension Network
8. https://postgres-slack.herokuapp.com/ - Slack channel
9. https://stackoverflow.com/questions/tagged/postgresql - StackOverflow
10. https://www.facebook.com/groups/postgres/ - FB (en)
11. https://www.facebook.com/groups/postgresql - FB (ru)
12. https://telegram.me/pgsql - Telegram (ru)
13. https://postgis.net/ - PostGIS
14. http://postgresql.org.il/ - PostgreSQL Israel
    http://pgday.org.il/ - PGDay Israel (March 19, 2018)

PostgreSQL Universal Database