

LIGO-Virgo Collaboration

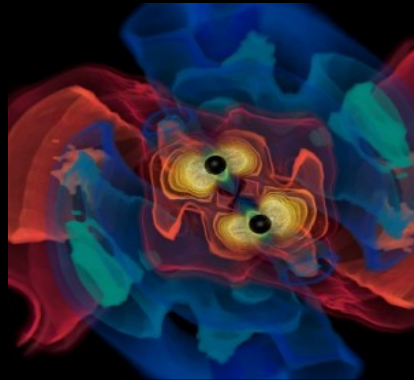
Frank Würthwein
OSG Executive Director
UCSD/SDSC

Why I am giving this talk

- I am not a member of either LIGO or Virgo.
- Unfortunately, we chose the date of HOW 2019 to coincide with the LIGO-Virgo Collaboration Meeting.
- Am thus giving this talk on behalf of Peter Couvares who helped me put these slides together.
 - All useful information is from him.
 - All mistakes are mine.

The Science

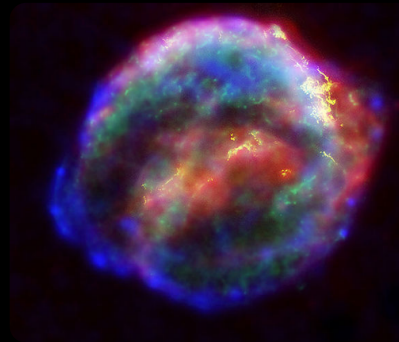
Gravitational Wave Astrophysics



Coalescing Binary Systems

Neutron Stars,
Black Holes

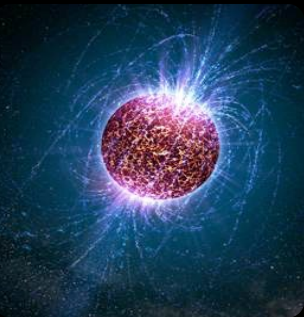
Credit: AEI, CCT, LSU



'Bursts'

Core collapse supernovae
Cosmic strings
Unknown

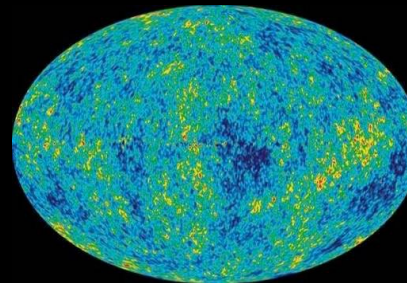
Credit: Chandra X-ray Observatory



Continuous Sources

Spinning neutron stars
crustal deformations,
accretion

Casey Reed, Penn State



Stochastic GW background
cosmological or astrophysical

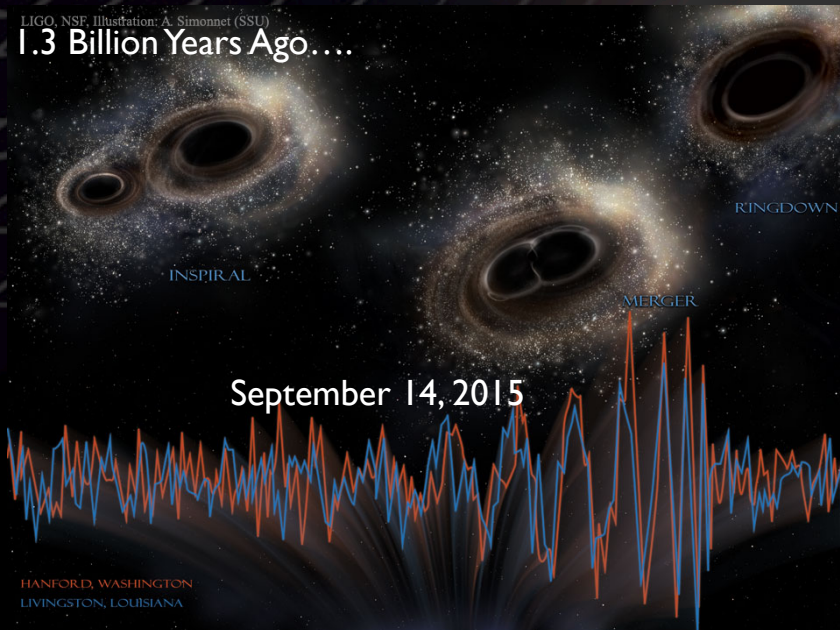
NASA/WMAP Science Team

Broadly speaking, 4 different types of wave phenomena are searched for.

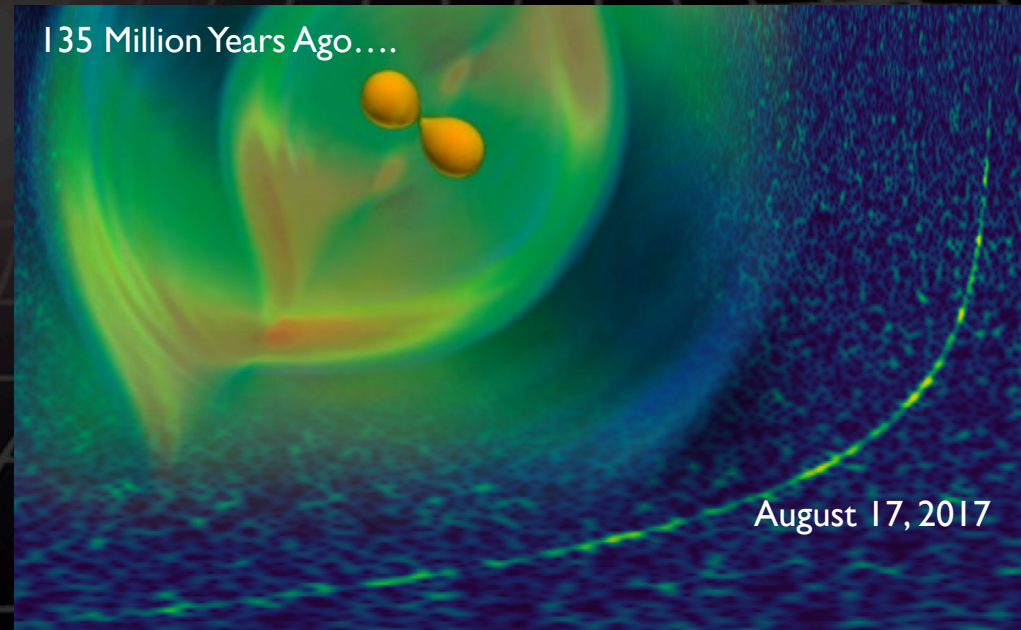


Coalescing Binary Detections

GW150914 and GW170817:
Two ground-breaking discoveries that opened a
new era in Gravitational Wave Astronomy



Binary Black Hole Coalescence



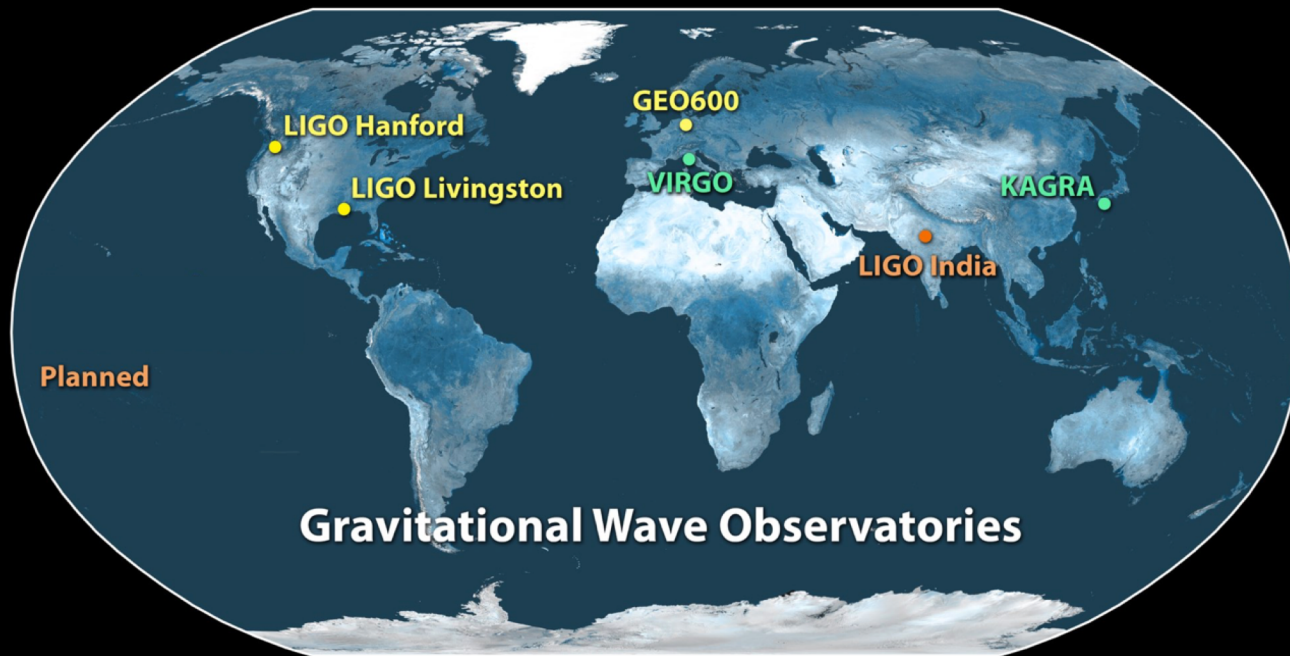
Binary Neutron Star Coalescence

Both come from just one of the 4 signatures that are being searched for.



Open Science Grid

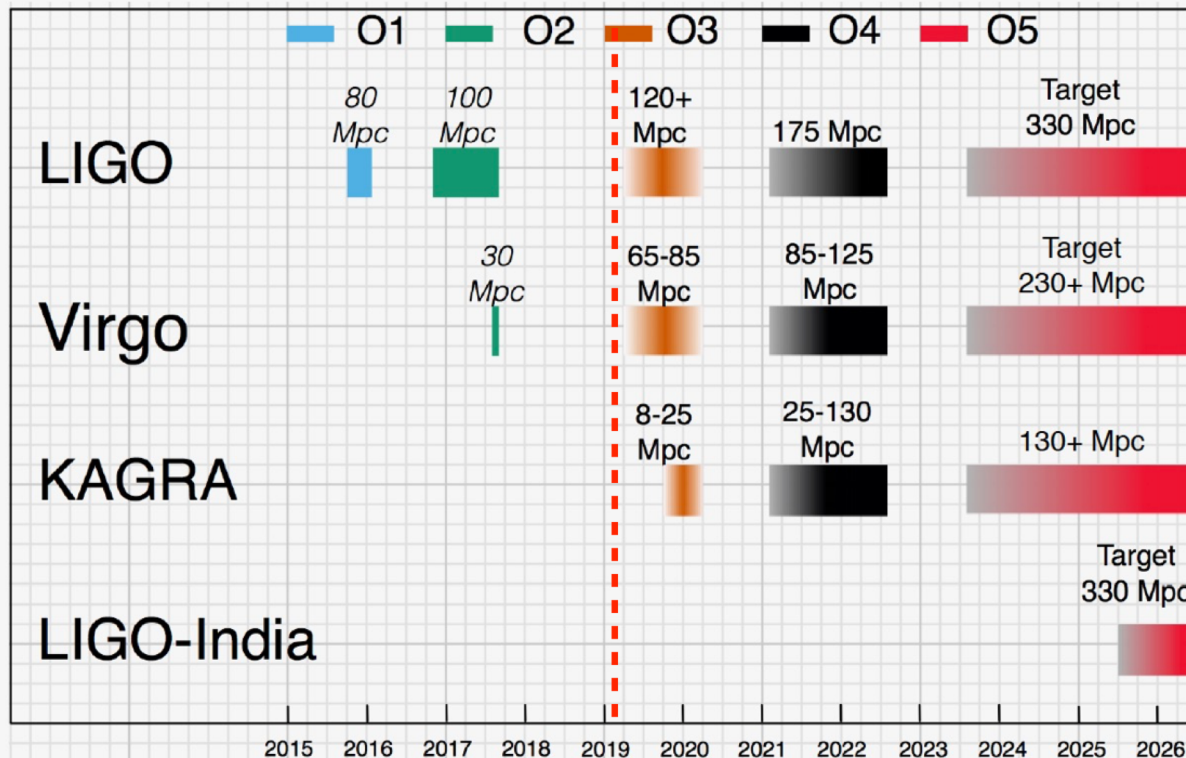
Multiple Instruments Working Together



A Global Quest



Near Term Observing Plan

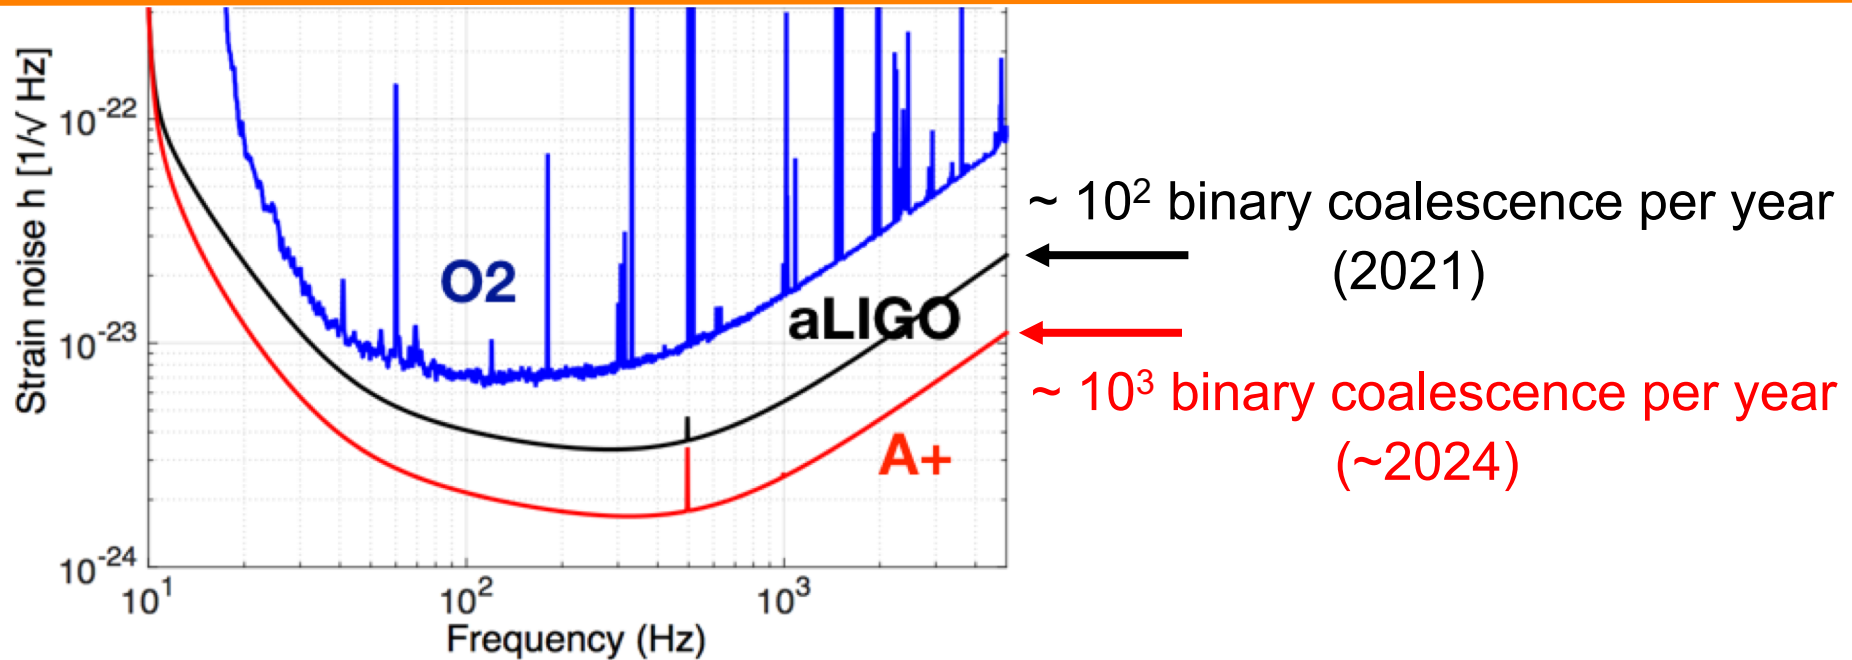


Prospects for Observing and Localizing Gravitational-Wave Transients with Advanced LIGO and Advanced Virgo and KAGRA —
<https://dcc.ligo.org/LIGO-P1200087/public>

Number of detected compact binary mergers $\sim R^3$



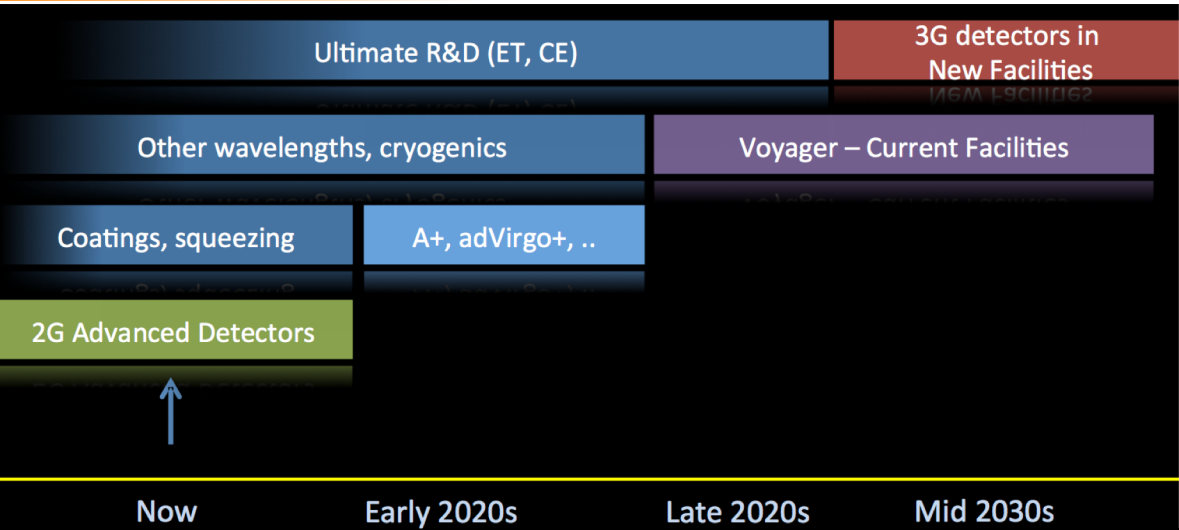
Evolution of Strain Sensitivity



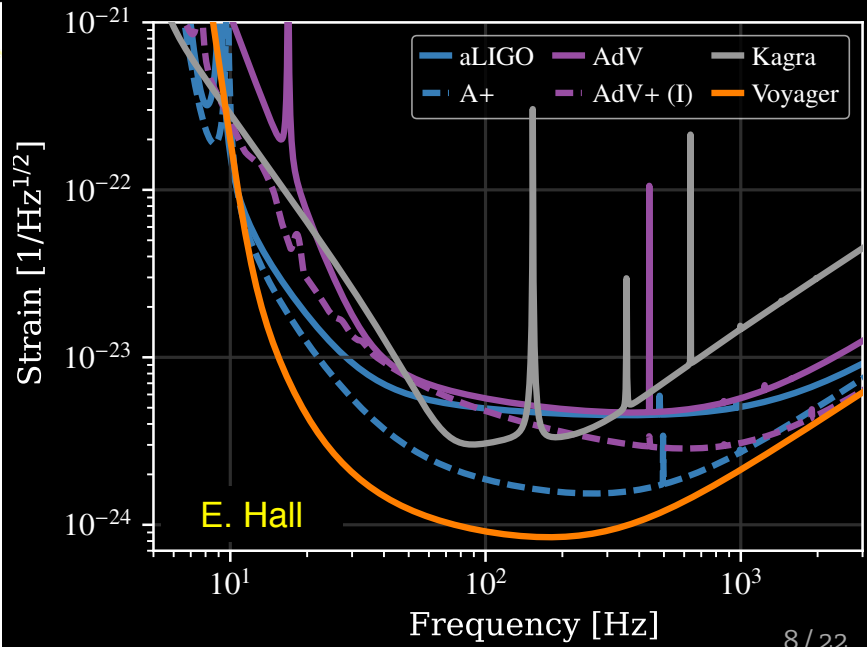
Rapid Reduction of Strain Noise within the next few years leads to massive increase in detections because of R^3 .



Longer Term Roadmap



Voyager ~ 10^4 binary coalescence/year





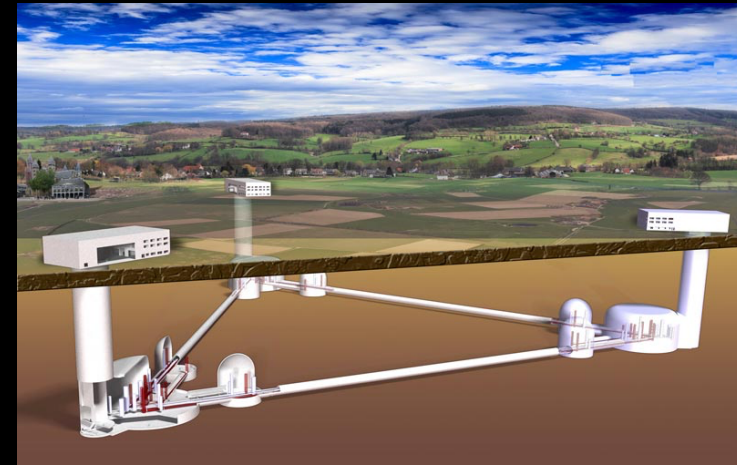
Multiple 3rd Generation Approaches

The 3rd Generation

$\sim 10^5$ binary coalescences per year (2030s)

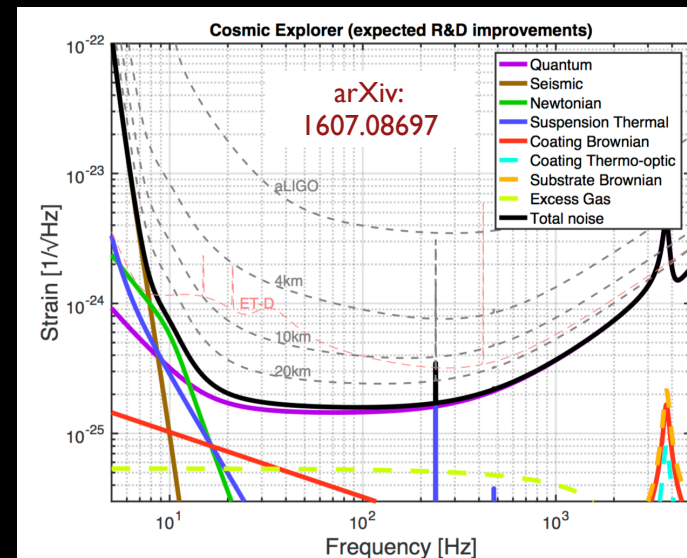
Einstein Telescope

- European conceptual design study
- Multiple instruments in xylophone configuration
- underground to reduce newtonian background
- 10 km arm length, in triangle.
- Assumes 10-15 year technology development.



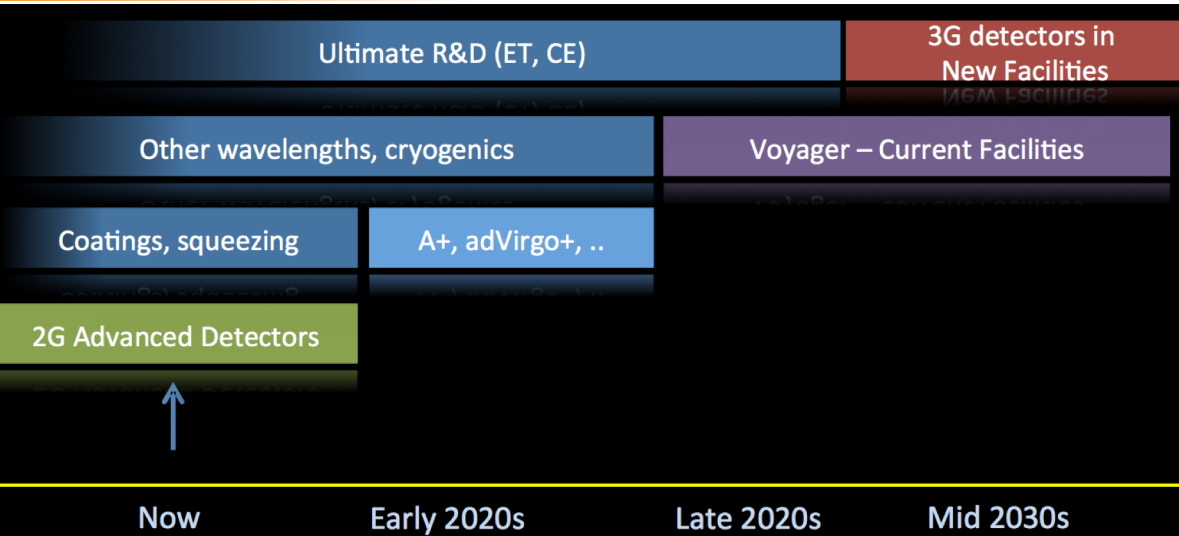
Cosmic Explorer

- NSF-funded US conceptual design study starting now
- 40km surface Observatory baseline
- Signal grows with length – not most noise sources
- Thermal noise, radiation pressure, seismic, Newtonian unchanged; coating thermal noise improves faster than linearly with length



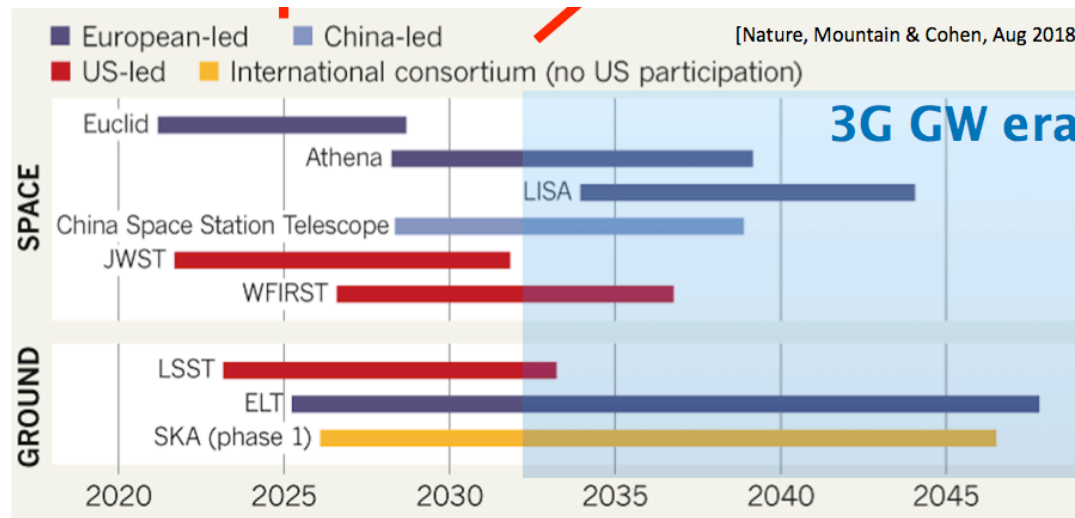


Multi-messenger Astrophysics



Many powerful instruments online concurrently

Credit: L. Barsotti



Computing

Ligo-Virgo Computing

Compute:

- For the third observing run (O3), the LVC expects to need ~500 million CPU core-hours of data analysis for ~80 astrophysical searches, followup activities, and detector characterization.
- The 10 most demanding of these 80 analyses comprise ~90% of the demand.
- Most of this computing is “pleasingly parallel” High Throughput Computing (HTC) for “deep” offline searches; ~10% is low-latency data analysis needed to generate rapid alerts (OPA)
- Currently ~90% provided by dedicated LIGO-Virgo clusters vs. ~10% from external shared computing resources — but growth of the dedicated resources has flattened while shared component is growing.
- Growing shared, external computing resources are presenting new distributed computing and data access challenges. GPU has been the most successful and cost-effective: deploying at scale for the first time in O3.
- Currently no cloud usage; no major technical obstacles, but logistics are unclear.

Data:

- LIGO $h(t)$ strain data is $O(10\text{TB})$ per IFO per observing year.
- LIGO raw data (all channels, full sample rate) is $O(1\text{PB})$ per IFO per observing year.
- No longer “big data” by 2018 standards — but non-trivial in a distributed HTC environment.

Commitment to Open Data

	2019												2020												2021									
	1	2	3	4	5	6	7	8	9	10	11	12	1	2	3	4	5	6	7	8	9	10	11	12	1	2	3	4	5	6	7	8	9	10
O1 Run																																		
GW150904																																		
GW151226+LVT151012																																		
O2 Run																																		
GW170104																																		
GW170814 + GW170817																																		
GW170608																																		
O3 Run (2 chunks)																																		

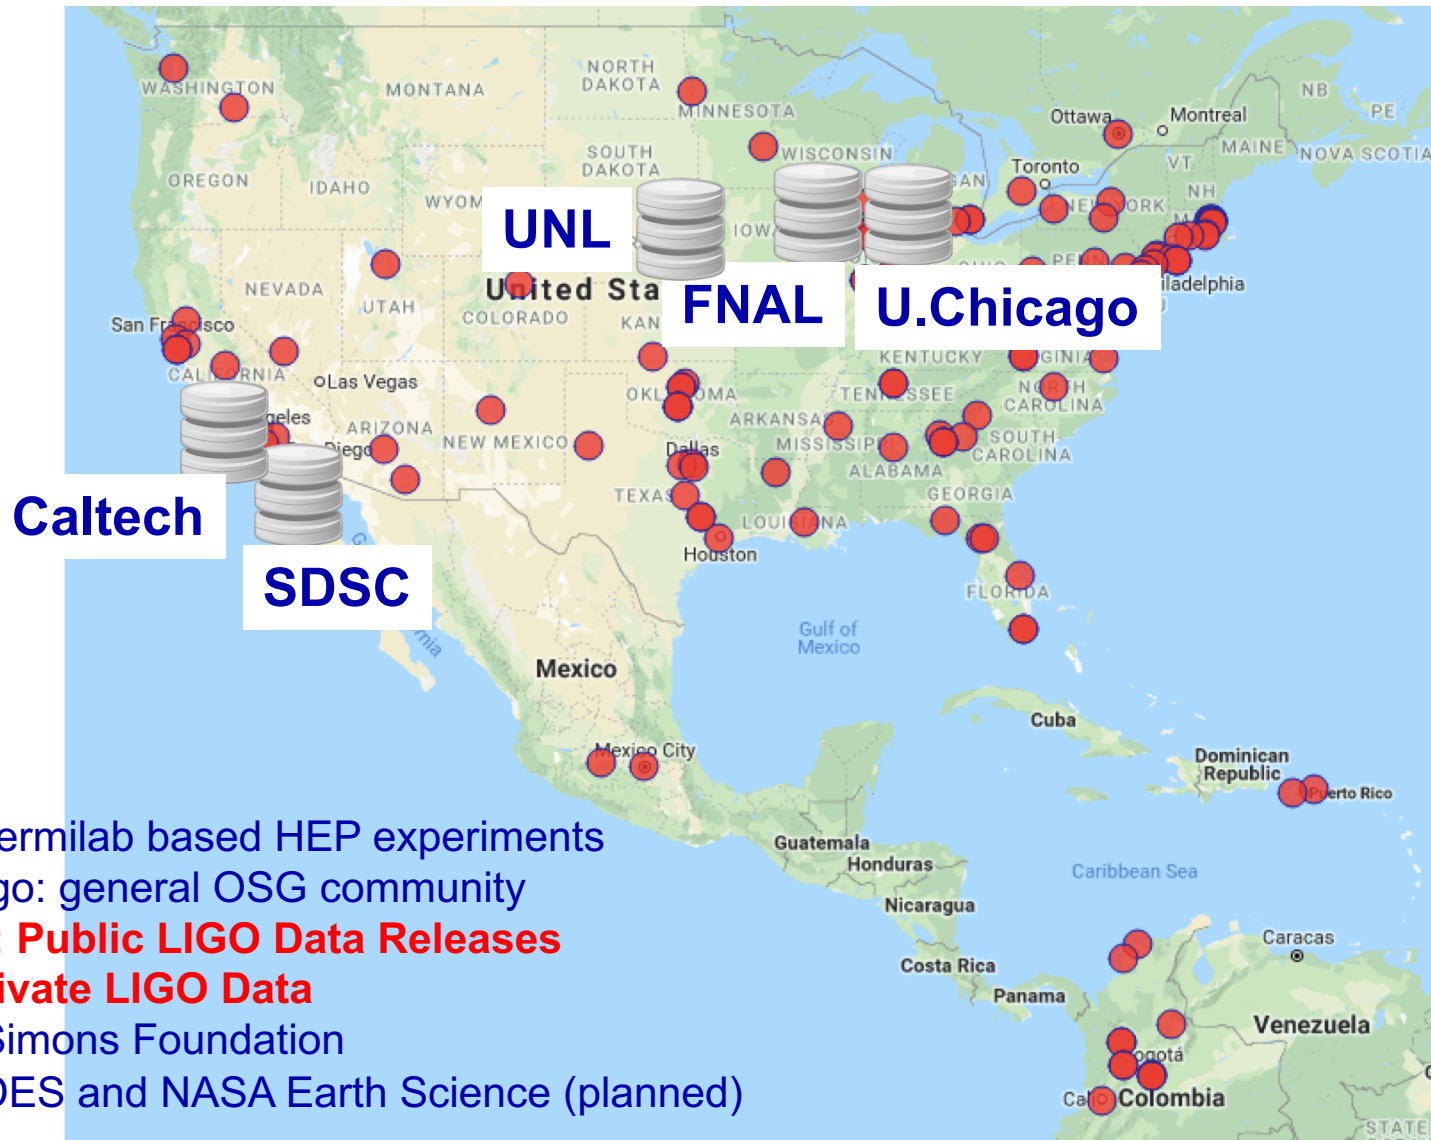
	Data Acquisition
	1.5 year proprietary period (as specified in the LIGO Data Management Plan)
	Open data

**See talk by Duncan Brown on use of this Open Data.
(Wednesday 15:00 OSG-AHM Session)**

Also, starting with O3: Open Public Alerts.
<https://www.ligo.org/scientists/GWEMalerts.php>



LIGO-Virgo use the OSG Data Federation based on Xrootd/cvmfs



FNAL: Fermilab based HEP experiments

U.Chicago: general OSG community

Caltech: Public LIGO Data Releases

UNL: Private LIGO Data

SDSC: Simons Foundation

NCSA: DES and NASA Earth Science (planned)



Caches in network & at endpoints

9 caches, one of which in EU

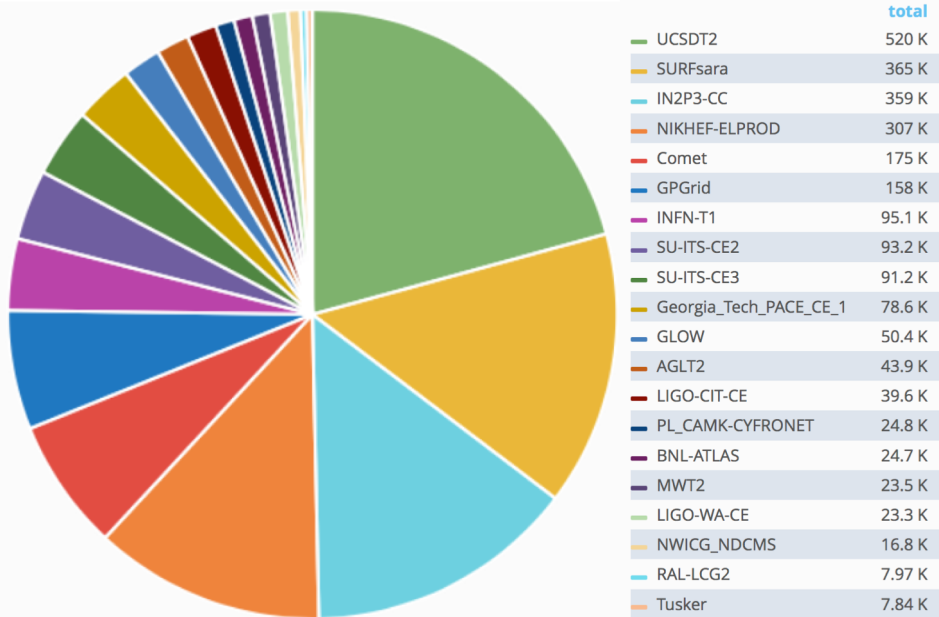


Last 30 days

	Directory ⚡	Bytes ⚡
GW public Data O1	/pnfs/fnal.gov/usr/dune	390.114TB
	/gwdata/O1	180.208TB
	/pnfs/fnal.gov/usr/minerva	99.352TB
	/pnfs/fnal.gov/usr/des	63.204TB
GW private Data	/user/ligo	58.211TB
	
GW public Data O2	/gwdata/O2	3.163TB

GW O1 Data second most popular in last 30 days
GW O2 Data only recently become public.

A case for caching



Top 20 core-hour contributors to LIGO-Virgo via OSG within last year.

**1/3 of it coming from WLCG T1 in EU
SURFSara, IN2P3, NIKHEF, INFN, RAL**

The **LIGO workflow** reuses each file $O(100)$ times.

Total data is only few TB but we moved many petabytes worth out of UNL before LIGO started using the caches in OSG.

Cache in Amsterdam is effective way to reduce transatlantic network traffic.

- As the gravitational wave community has become a global enterprise using WLCG resources via OSG and EGI middleware we need better coordination:
 - Issues with recognition of certificate authorities
 - CA bundles from OSG and EGI are not identical. The CA that is used by LIGO users is not in EGI bundle.
 - Issues with network routing from WLCG to cache at University of Amsterdam
 - E.g. IN2P3 – NYC – University of Amsterdam
- Concerned of hitting other problems in the future unless we make an effort to include LIGO-Virgo needs in existing coordination via WLCG.

Summary & Conclusion

- **LIGO-Virgo Collaboration has rapidly expanding science program on WLCG resources.**
 - Supporting private and public data.
- Roughly x1000 increase in coalescent binary detections/year expected within the next couple decades.
- Additional interferometers coming online.
- Desire for existing coordination between OSG & EGI within WLCG to include focus on GW community needs.