

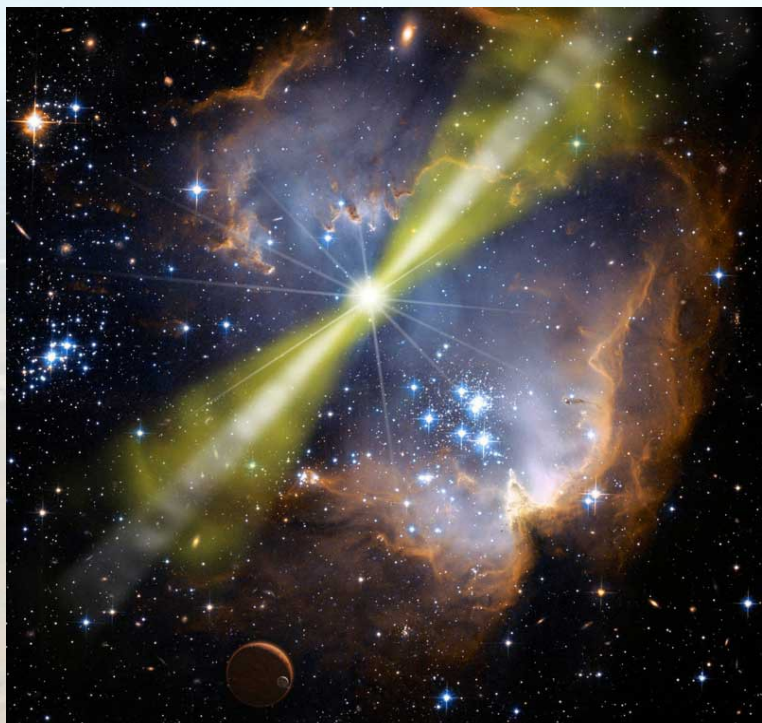
Computing issues in experimental gravitational wave research: Virgo computing

Gergely Debreczeni

Wigner RCP
Virgo Computing Coordinator

(Debreczeni.Gergely@wigner.mta.hu)

Content



Artists's view of a gamma-ray burst

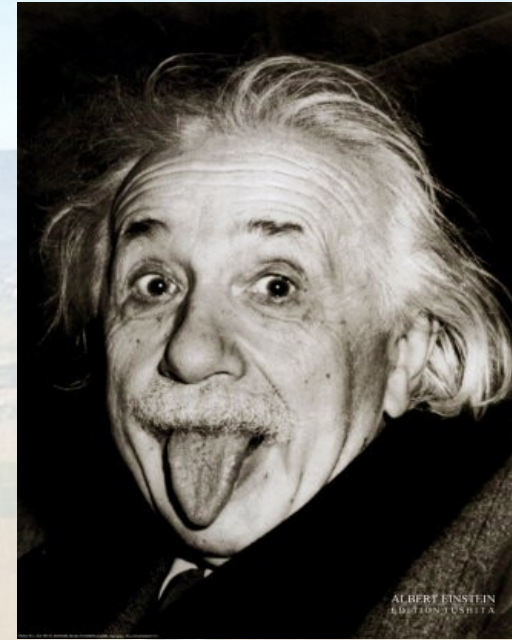
Credit: NASA/Swift/Mary Pat Hrybyk-Keith and John Jones

(... a possible source of gravitational waves)

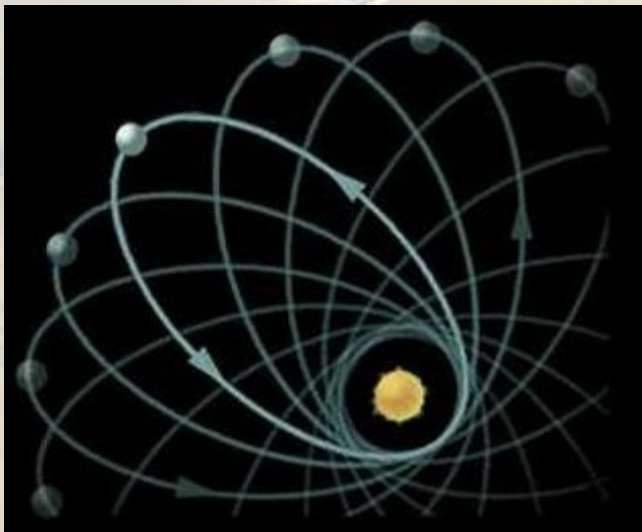
- **Gravitational waves**
 - What are to so called gravitational waves ?
- **Virgo experiment**
 - A one slide introduction to the Virgo experiment
- **Computing model walkthrough**
 - Scale, type and strategies in Virgo computing
- **Data transfer**
 - Transfer of measurement data to CCs
- **Analysis types**
 - Computing needs of anaalysis workflows
- **Pipeline execution problems**
 - Problems related to pipeline execution
- **GPU perspectives**
 - How GPUs can help us in cheaper computing
- **Cloud solution**
 - Why a computing cloud would be excellent for Virgo
- **Future plans**
 - Things we would like to work on in the future

About gravitational waves

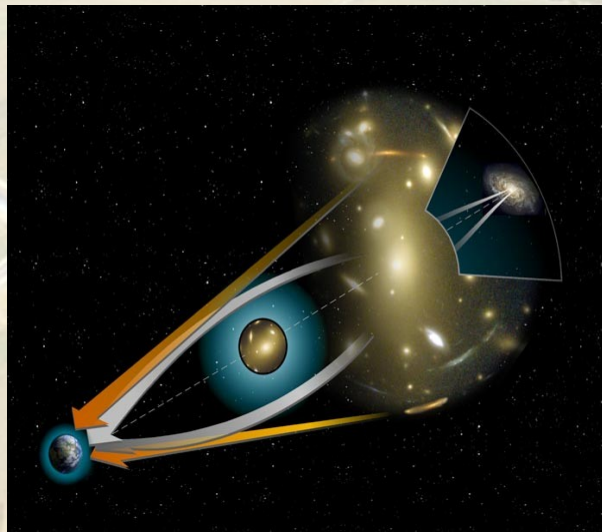
- Several predictions of Einstein's general relativity has already been confirmed experimentally, like
 - advance of perihelium
 - gravitational lensing
 - change of elapse of time in strong gravitational fields, etc...
- but we have only indirect (but very good) proof for the existence of gravitational waves



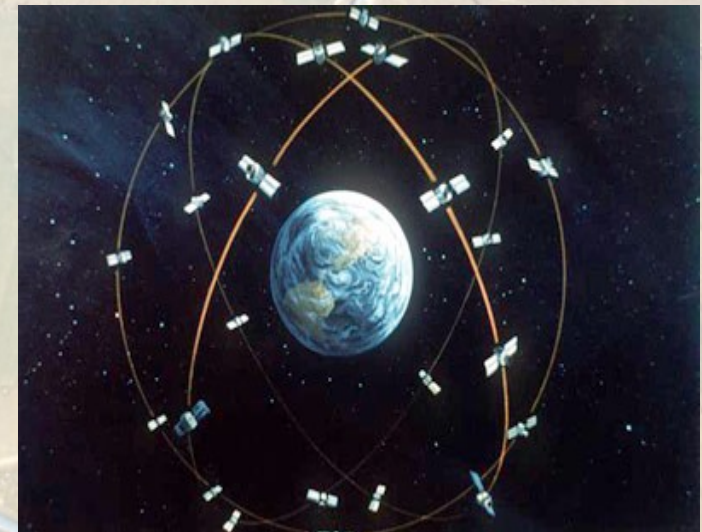
The direct detection of GWs are the goal of the VIRGO experiment.



2015 may 20. - GPU Day

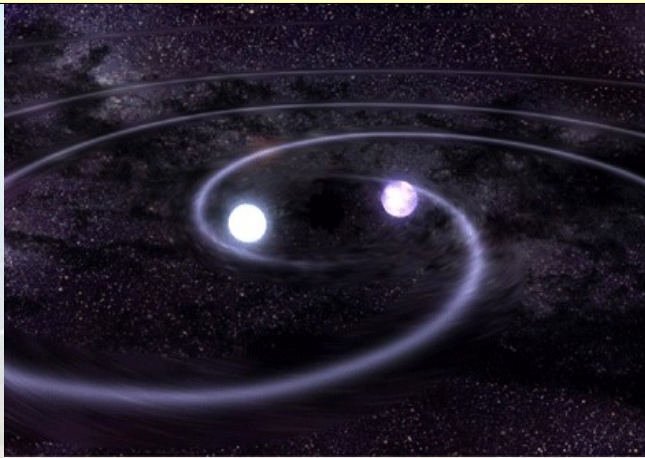


Gergely, DEBRECZENI - GRB forecastin

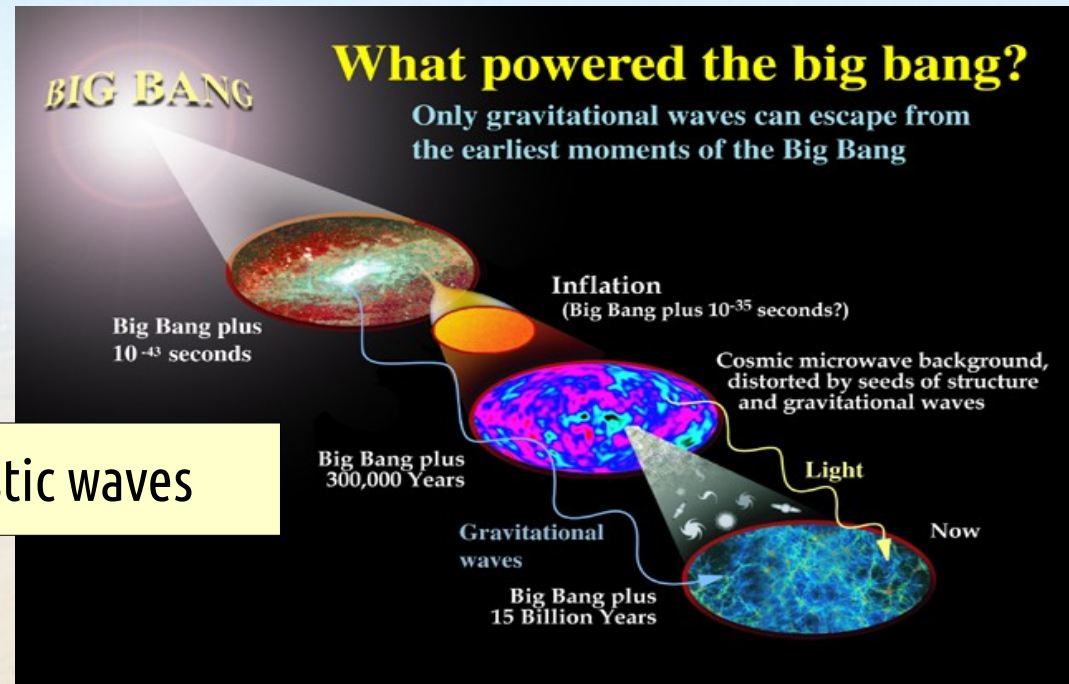


Classification of GW sources

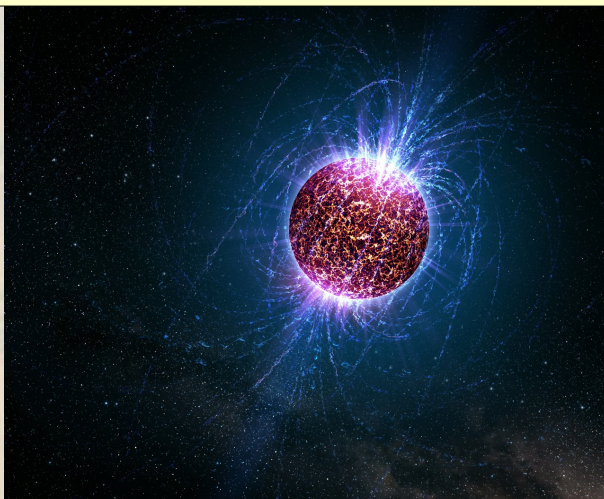
Compact Binary Coalescence



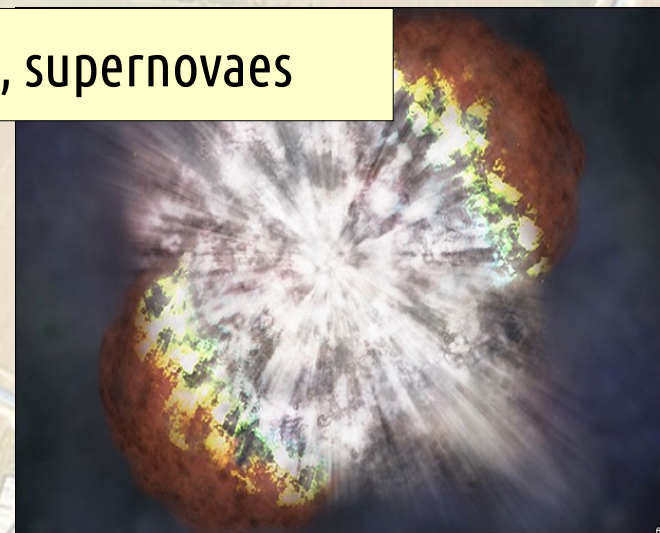
Stochastic waves



Pulsars, neutron stars

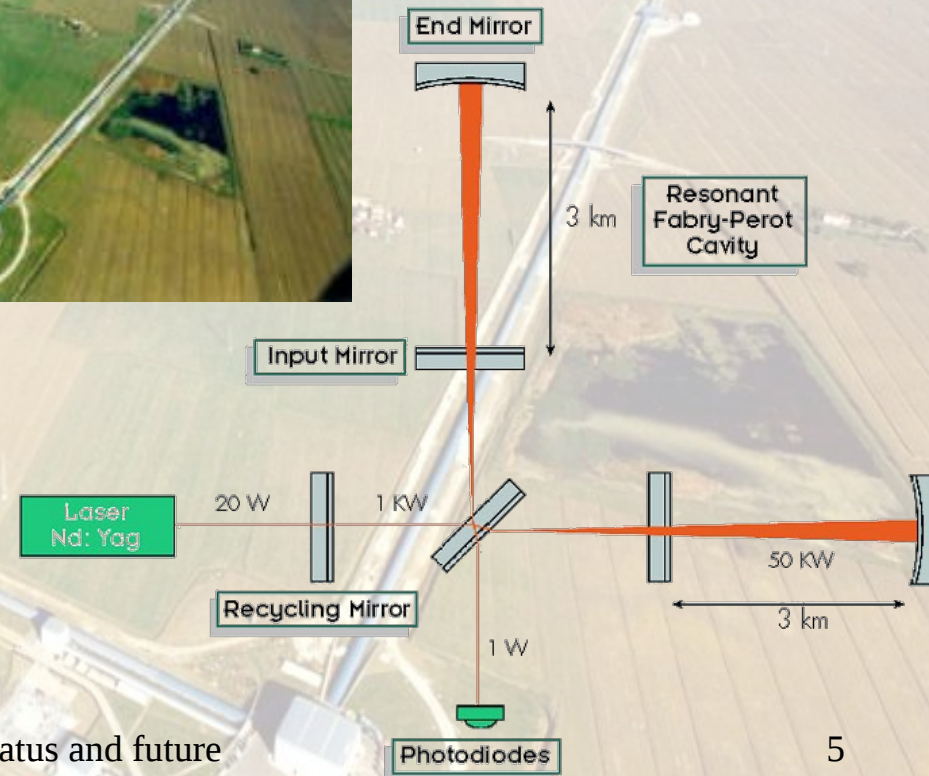
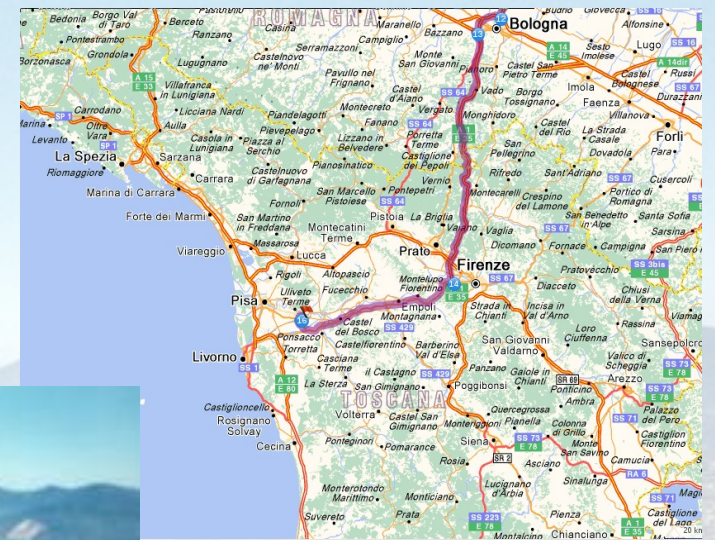


Bursts, supernovae



The Virgo experiment

- The Virgo detector is located in the site of the European Gravitational Observatory (EGO) in Cascina, near Pisa, Italy.
- Construction finished in 2003
- It is now a european collaboration including France, Italy, Hungary, Netherland, Poland
- Working together with LIGO (Laser Interferometer Gravitational-wave Observatory), synchronized observations and coordinated analysis
- So far, approximately c.c 20 month of data taking
- Currently under upgrade, will start to collect scientific data in early 2016



The Virgo Computing Model in 1 slide

- 1) Data is taken by the 3 detector (Virgo + 2 LIGO detector)
- 2) **Online** (low-latency) analysis happens on the measurement site.
- 3) Data is stored at site temporarily in a circular buffer (typically for 6 month)
- 4) All data is transferred with c.c. 1 day latency to external CCs and stored in 3 different location (Lyon, CNAF, Ligo site)
- 5) For **offline** analysis Virgo uses its two main CCs and the INFN Grid sites, while LIGO uses dedicated Condor clusters and XEDE supercomputer resources
- 6) Offline analysis based on Condor DAG, Pegasus DAX and shared file system based workflows and/or simple EMI Grid submission mechanisms.
- 7) No collaboration-wide unified job scheduler is used. Multiple solutions.
- 8) Analysis code must undergo serious review, coordinated prioritisation and optimisation efforts. Reviewed code is tagged, freezed.

Type of measurement data

- The measurement data is a **simple time series**, sampled at 20 kHz, then downsampled to 16 kHz and 4 kHz.
- There are hundreds of auxiliary environmental channels, some of them with much lower frequencies and sampling rate.
- Amount of data is a **few hundred TB / yr** but its **arithmetic complexity is much higher** than that of the HEP experiments.
- Depending on the source to be examined / discovered many different kind of analysis is crunching this data with computing requirements differing by order of magnitude.
- Different analysis requiring different input data size and computing architectures
- Gravitational wave analysis is **compute intensive** not data intensive

Data transfer

- Measurement data produced in Cascina on the Virgo site.
- **Low latency** online analysis happens in place and data is stored temporarily on site using a circular buffer of length of several month.
- For **offline analysis** this data must be transferred to the computer centers.
- For this purpose a transfer tool was developed by EGO Computing department. Existing out-of-the box solutions such as FTS were too heavy to maintain.
- It uses the
 - lcg-tools to transfer data to CNAF
 - and the iRod client installation for IN2P3
- The amount of data is not overwhelming, in principle data transfer should not be a challenge
- Data from LIGO detectors does also get copied to our CCs. This data consist(ed) of small file which must be merged to bigger ones in order to fit better for HPSS storage. Now this problem had been solved

Data transfer issues

- Virgo is using different data transfer tools for different data transfers. Each of them is easy and works reliably, however the need for using multiple backend needs extra manpower and development.
- A preferred solution would be to use the LDR (LIGO Data Replicator) data transfer framework all over in the LIGO-Virgo collaboration which can communicate with legacy GSIFTP backends and perform reliable data transfer.
- GSIFTP is not available in IN2P3 which is a problem, probably will be solved soon

Data file and metadata catalogs

- Some analysis is using the LFC (Lightweight File Catalog). Its use is easy and we found no problem with it.
- The Diskcache software is used to catalog files available in a given CC and respond to queries of the pipelines with physical location of files.
- Typical query includes GPS times, detector name, channel name -> there is not too much metadata that can be assigned to measurement data.
- No file metadata catalog is used so far, there is not too much need for that. Many information is in file names and can be easily queried.

Analysis pipelines

- **CW** (Continuous waves) - Rotating, asymmetric neutron stars
 - The most compute intense pipeline, practically can consume all available resources. Sensitivity goes like $1/\sqrt{T}$, where T is the duration of data chunk in question. As a result one must restrict parameter space explored -> **computing constraints yields scientific limitation.**
- **CBC** (Compact Binary Coalescence) - Gravitational waves emitted by coalescing binary neutron stars or binary black holes
 - Very compute intensive, theoretical templates are tested against the measurement data by means of matched filter. With the decrease of low frequency cutoff compute costs grows exponentially.
- **Burst** - Explosions, supernovas, unmodelled transient sources
 - Very similar but more generic than CBC. Sensitivity is comparable (~c.c 30% less).
- **Stochastic** - Search for stochastic gravitational waves of galactic or primordial origin
 - Important from physics point of view, but has negligible compute cost.

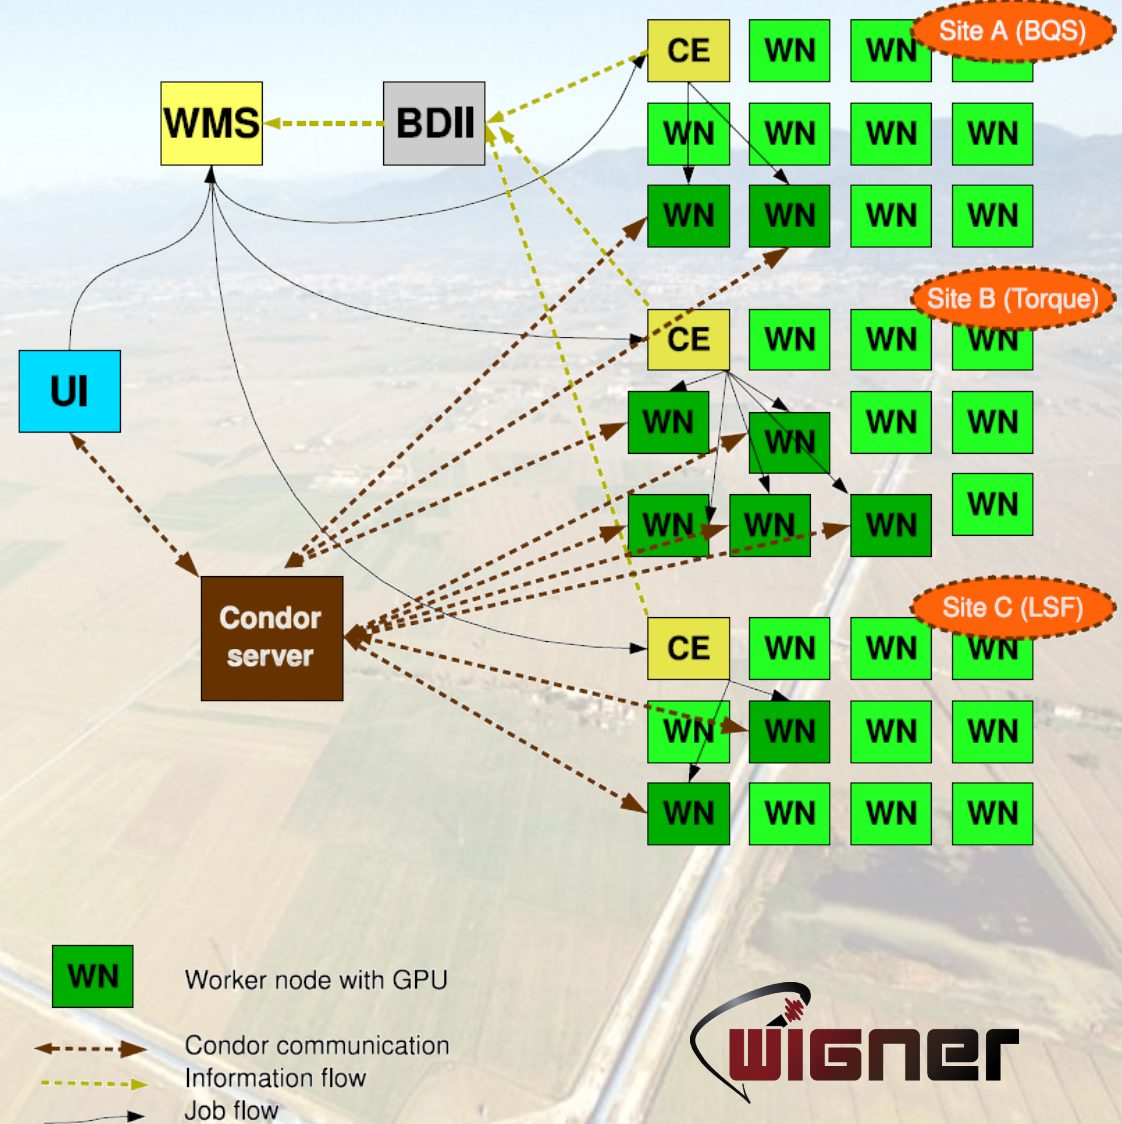
Analysis pipeline problems

- LIGO and Virgo collaboration is working closely together.
- However LIGO collaboration is larger by a factor of $> \times 10$
- As a result many important pipeline development is dominated by LIGO colleagues and are tailored to LIGO resources
- As a consequence those pipelines cannot be executed out of the box on our resources, but requires quite some porting effort.
- This effort is not a one time action, but needs continuous attention -> very expensive in terms of time and manpower...that Virgo cannot afford.
- Many attempts have been made to overcome this difficulty including
 - the set up of a pilot pool framework
 - using the Pegasus scheduler
 - examining the possibility of using the Dirac jobmission framework
 - thinking on virtualized Condor cluster, i.e. a Cloud

The Virgo Pilot Pool - I

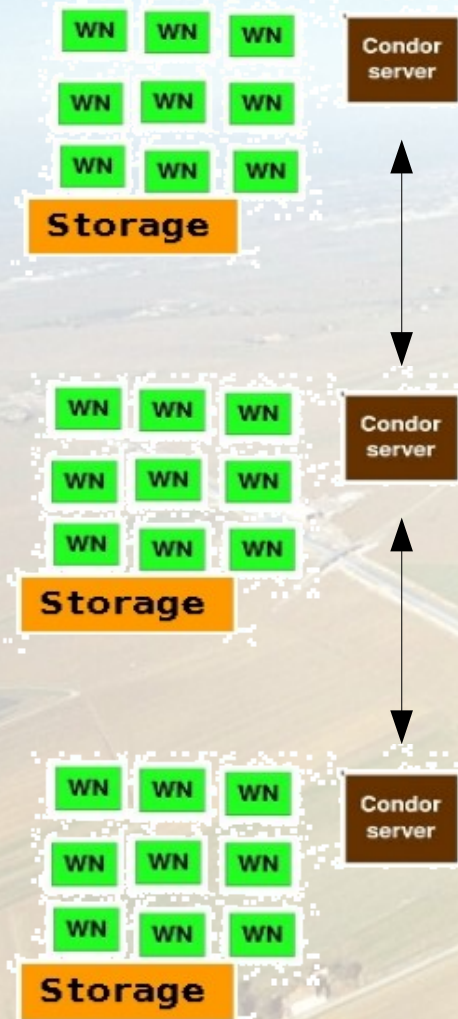
The Virgo Pilot Pool properties:

- Homogen infrastructure over the inhomogen Grid
- Less administrative interaction/delay
- User transparent mechanisms
- Low latency submission
- Global priorities
- Late-binding to resources
- No stuck-in jobs
- Improved job failure rate due to pilot prechecks
- Interactive login
- Smooth interaction interoperability with LDG/OSG.

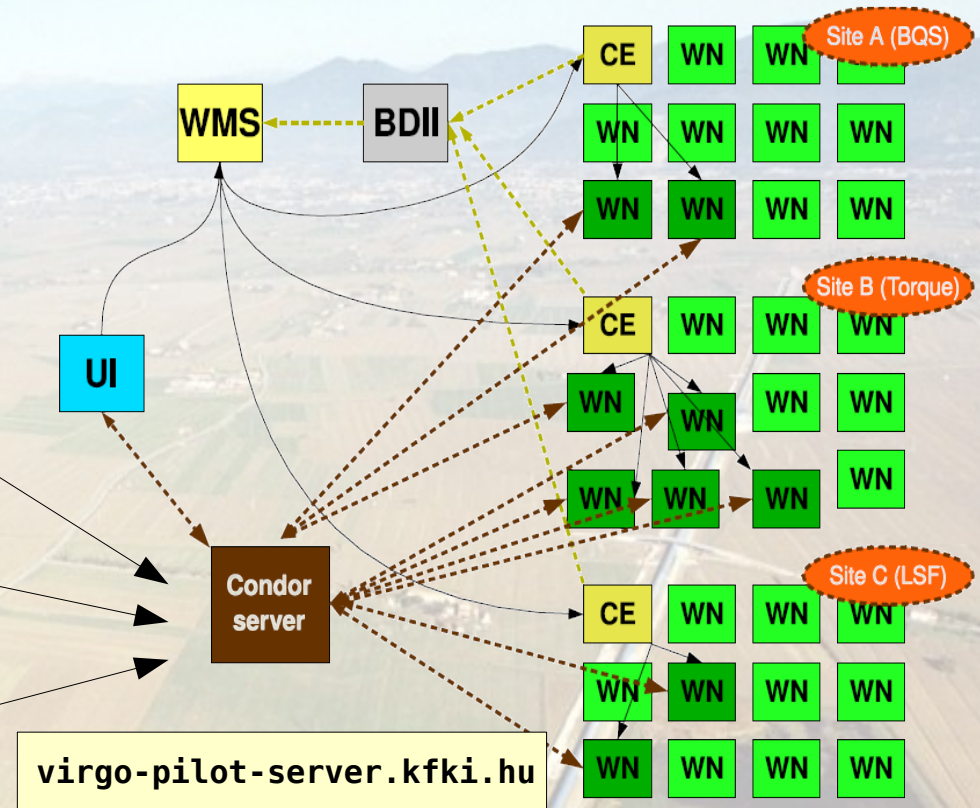


The Virgo Pilot Pool - II

Ligo clusters



Virgo EGI sites



Worker node with GPU



Condor communication



Information flow

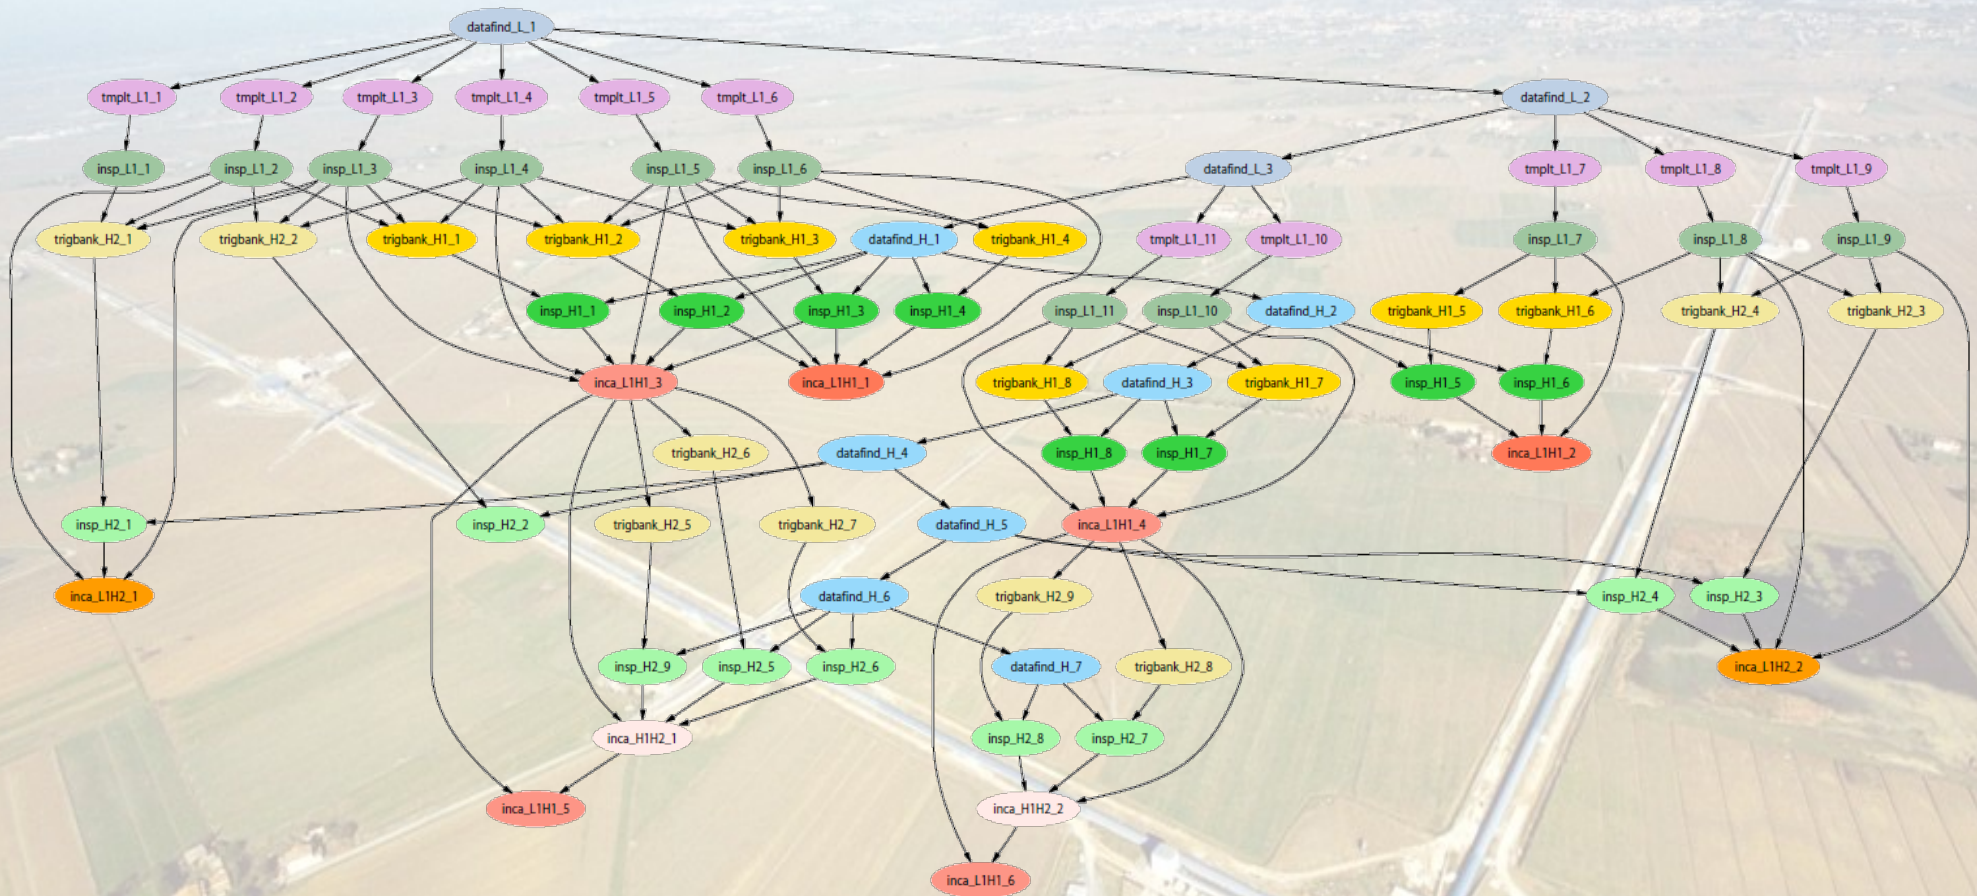


Job flow



The Virgo Pilot Pool - III

- Mapping of abstract workflows like DAGs/DAXes to the Grid is now easily possible with the Virgo Pilot server.
- Complex and relational workflow handling is/was missing from Cream/WMS.



Cloud 4 Virgo

- Out of the above possibilities each of them has some serious drawbacks except the Cloud solution
- Having an OpenStack based virtual Cloud Condor installation at CNAF would solve almost all our problem including,
 - pipeline porting
 - architectural difference respect to LIGO clusters
 - training of people and the must to learn multiple submission mechanisms
 - real sharing of our resources
 - ease of GPU access and GPU - CPU matching, allocation problems
 - better monitoring of user jobs

GPUs for analysis

- Many search algorithm can be accelerated by making use of operation level parallelizability offered by various many-core hardware such as GPUs. Such examples are:
 - FFT, vector operations, reduce, max finding, clustering in **CBC analysis pipelines**
 - FFT, 2D thresholding, differential Hough map creation, integration, peak finding, etc.. in **CW analysis**
- There are multiple tool developed to allow easier use of GPUs by less advanced programmers, such as:
 - GWTools - An OpenCL based templates C++ generic algorithm library for GW searches
 - pyCBC - CUDA based set of Python algorithm used in CBC analysis
 - CB - Compute Backend - offers a unified host code for CUDA and OpenCL, so there is no need to write the code twice for NVidia and AMD cards
- GPUs will play crucial role in the following years probably even for the discovery
- Typical full-pipeline **accelerations** experienced are ranging from **x30 to x120**

Optimisation, prioritisation

- The LIGO - Virgo analysis software stack went through on a serious process of review, benchmarking and optimisation.
- The process was triggered by the NSF review of LIGO request for XEDE resources
- It has a very positive effect on the quality, organisation and performance of the code used for analysis in the Collaborations.
- Analysis type based compute resource request estimation, logging and accounting and prioritisation is just under introduction in the collaborations
- A common unit of measure for called „Service Unit” SU has been introduced, since we observed that HS06 numbers are not alway accurate enough in reflecting the ratios of performances of a specific CPU cores for specific analysis. 1 SU corresponds to the performance of an Intel Xeon E5-2650 core.
- LVCN - the LIGO - Virgo Computing Network was established

LVCN supply table

LI60-Virgo O1 CPU/GPU/MIC Inventory (March 2015)																																								
DEPLOYED SYSTEMS					INDIVIDUAL DEVICE STATISTICS										SYSTEM TOTAL					RELATIVE PERF					AVAILABLE CORES (E5-2670 equivalent)															
Tag	SITE	Org	Arch	Model	Qty	Peak GLOPS	SPECfp GLOPS	FFT GLOPS	pycbc (4k)	pycbc (2k)	gstlal (2k)	Core	Clock MHz	Cache MB	max flop per cycle	Mem GB	Peak GLOPS	CPU core	GPU core	Mem GB	SPECfp %	Full Device Performance	FFT	pycbc(4k)	pycbc(2k)	gstlal(2k)	Dev %	Dev core	Prod %	Prod core	Peak core	O1 rank	O2 rank	O3 rank	Deployed date	EOL	Comment			
CIT	CIT	Lab	CPU	Opteron 275	662	18	15					2	2,200	2	4	2	11,651	1,324	1,324	0.06							100	331	0	0	0				02/2006	2015				
CIT			CPU	Opteron 2376	328	74	54						4	2,300	6	8	8	24,141	1,312	2,624	0.23							100	590	0	0	0	1			03/2009				
CIT			CPU	Xeon E5-2670	96	333	240	102.9	43,875	97,488	44,000	8	2,600	20	16	64	31,949	768	6,144	1.00	1.00	1.00	1.00	1.00	1.00	1.00	0	100	768	768		1	1	1	12/2012	8/2013				
CIT			CPU	Xeon E3-1241v3	200	448	153		27,903	71,052			4	3,500	8	32	32	89,600	800	6,400	0.64			0.64	0.73			0	0	100	1,020	1,020		1	1	1	4/1/2015		aLIGOFY15 procureme	
CIT			GPU	Tesla K10	64	4,580		378					3,072	745			8	293,120	196,608	512		3.67						0		100						12/2012				
CIT	LLO	LHO	MIC	Xeon Phi 5110P	32	2,022						60	1,500			8	33,696	1,920	256							0		100							8/2013					
OBS			CPU	Opteron 2376	504	74	54						4	2,300	6	8	8	37,094	2,016	4,032	0.23							100	907	0	0	0		1			03/2009			
OBS			CPU	Xeon E3-1241v3	100	448	153						4	3,500	8	32	32	44,800	400	3,200	0.64							0	0	100	510	510		1	1	1	4/1/2015		aLIGOFY15 procureme	
OBS			CPU	Opteron 2376	504	74	54						4	2,300	6	8	8	37,094	2,016	4,032	0.23							100	907	0	0	0		1			03/2009			
OBS			CPU	Xeon E3-1241v3	100	448	153						4	3,500	8	32	32	44,800	400	3,200	0.64							0	0	100	510	510		1	1	1	4/1/2015		aLIGOFY15 procureme	
HAN	HAN	HAN	CPU	Xeon X3220	1,176	38	40					4	2,400	8	4	8	45,158	4,704	9,408	0.17							50	784	0	0	0		1			03/2008	2014/5			
HAN			CPU	Xeon E3-1220v3	1,820	397	137	40	28,511	64,397			4	3,100	8	32	16	722,176	7,280	29,120	0.57	0.39		0.65	0.66			0	0	50	4,156	8,311		1	1	1	02/2014			
HAN			CPU	Xeon E3-1231v3	270	435	151	55			37,600		4	3,400	8	32	16	117,504	1,080	4,320	0.63	0.54					0.85	0	0	50	680	1,359		1	1	1	02/2014		0	
HAN			CPU	Xeon E5-1650v2	52	336	220						6	3,500	12	16	64	17,472	312	3,328	0.92							0	0	50	191	381		1	1	1	02/2014			
HAN			CPU	Xeon E5504	62	64	57						4	2,000	4	8	12	3,968	248	744	0.24							0	0	50	59	118				01/2010	2015?			
HAN	SYR	LSC	GPU	Tesla C1060	116	933						240	1,300			4	108,228		27,840	464															01/2010	2015?	Atlas GPU numbers cor			
HAN			GPU	Tesla C2050	101	1,930							448	1,150			3	104,030		45,248	303														01/2010	2015?				
HAN			GPU	GT640	459	691							384	900			2	317,169		176,256	918															03/2014				
HAN			GPU	GTX750	270	1,044							512	1,080			2	281,880		138,240	540															06/2014				
SYR			CPU	Xeon X3220	80	38	40							4	2,400	8	4	8	3,072	320	640	0.17							50	53	0	0	0	1			10/2007	10/2015	Sugar: Confirm availab	
SYR	UWM	UWM	CPU	Xeon X5650	667	128	125		23,808	43,864		6	2,668	12	8	24	85,419	4,002	16,008	0.52			0.54	0.45			0	0	100	2,779	2,779		1	1	1	2015	2,018	Crush: MoU not signed		
SYR			GPU	GTX 580	172	1,581							512	722		2	271,949		88,064	258							50		0				1			10/2010	10/2015			
UWM			CPU	Xeon X5650	320	128	125						6	2,670	12	8	24	41,011	1,920	7,680	0.52							100	1,333	0	0	0	1			1/2010	2015			
UWM			CPU	Xeon E5-2650	184	256	210						8	2,000	20	16	32	47,104	1,472	5,888	0.88							0	0	90	1,159	1,288		1	1	1	9/1/2012	9/2017		
UWM			GPU	GTX 690	4	5,620							3,072	1,000			2	22,480		12,288	8							0		100							9/2012	9/2017		
UWM	IUCAA	Cardiff	GPU	GTX 580	4	1,581						512	722			2	6,324		2,048	6							0		100							9/2012	9/2017			
UWM			GPU	Tesla K10	10	4,580							3,072	745			8	45,800		30,720	80							0		100							9/2012			
UWM			GPU	Tesla M2075	10	1,030							448	1,150			6	10,300		4,480	60							0		100							9/2012	9/2017		
IUCAA			CPU	Xeon E5-2670	60	333	240						8	2,600	20	16	64	19,968		3,840	1.00							25	120	0	0	0		2	2	2	5/2013	8/2017		
Cardiff			CPU	Xeon X5660	128	134	125						6	2,800	12	8	48	17,203	768	6,144	0.52								0	0	67	357	533		1	1		12/2012	12/2016	
Cardiff	KGWG	CNAF	CPU	Xeon E5-2690v3	100	998	400					12	2,600	30	32	64	99,840	1,200	6,400	1.67							0	0	67	893	1,333		1	1	1	01/2015				
KGWG			CPU	Xeon E5645	30	115	106						6	2,400	12	8	12	3,456		360	0.44							50	53	0	0	0				01/2012	2016?			
KGWG			CPU	Xeon E5645	40	115	106						6	2,400	12	8	24	4,608	240	960	0.44							50	71	0	0	0				01/2012	2016?			
KGWG			CPU	Xeon X5650	4	128	125						6	2,670	6	8	24	513	24	96	0.52							50	8	0	0	0				01/2010	2016			
KGWG			CPU	Xeon X5650	18	128	125						6	2,670	6	8	12	2,307	108	216	0.52							50	38	0	0	0				01/2010	2016			
CNAF	Wigner	Nikhef	CPU	Heterogeneous	125	166	120					8	2,600	8	64	20,800	1,000	8,000	0.50								0	0	80	400	500		1			01/2011	8/2013	Remaining 20% availab		
Lyon			CPU	Xeon E5645	50	115	210						6	2,400	12	8	96	5,760	300	4,800	0.88							0	0	100	350	350		1			01/2013	2018		
Rome			CPU	Xeon E5410	52	154	70						8	2,400	12	8	8	7,987	416	416	0.29							0	0	80	97	121		1			01/2008			
Wigner			GPU	Opteron 6376	12	147	200						16	2,300	16	4	32	1,766	192	192	384	0.83							0	0	90	72	80		1			1/2013		
Wigner			GPU	AMD 7970	16	4,301							2,048	1,000			3	68,816		48									0		90						1/2013			
Nikhef	Polgraw	Polgraw	CPU	Xeon E5520	45	72	150					4	2,260	8	8	24	3,254	180	1,080	0.63							0	0	100	225	225		1					?		
Polgraw			CPU	Xeon E5-2665	16	307	235						8	2,400	20	16	64	4,915	128	1,024	0.98							0	0	100	125	125		2			1/2015	2018	To be used for Virgo CW	
Polgraw			CPU	Opteron 6172	20	101	150						12	2,100	12	4	32	2,016	240	640	0.93							0	0	100	100	100		2			1/2012	2016	To be used for Virgo CW	
Polgraw			CPU	Xeon E5645	80	115	106						6	2,400	12	8	24	9,216	480	480	1,920	0.44							0	0	100	283	283		2			?	2016	To be used for Virgo CW
Polgraw			GPU	GTX Titan	4	4,709							2,688	1,000			24	18,836		10,752	96							0		100							1/1/2014	?	To be used for Virgo CW	
EGO	EGO		CPU	Xeon E5-2650	13	256	210					8	2,000	20	16	128	3,328	104	1,664	0.88							0	0	100	91	91		1			1/1/2015	9/2017			
Tag	SITE	Org	Arch	Model	Qty	Peak GLOPS	SPECfp GLOPS	FFT GLOPS	pycbc (4k)	pycbc (2k)	gstlal (2k)	Core	Clock MHz	Cache MB	max flop per cycle	Mem GB	Peak GLOPS	CPU core	GPU core	Mem GB	SPECfp %	Full Device Performance	FFT	pycbc(4k)	pycbc(2k)	gstlal(2k)	Dev %	Dev core	Prod %	Prod core	Peak core	O1 rank	O2 rank	O3 rank	Deployed date	EOL	Comment			

An example for demand - supply - benchmark matchmaking sheet

2015 may. 25. - INFN CNAF

Gergely, DEBRECZENI - Status and future of gravitational wave computing

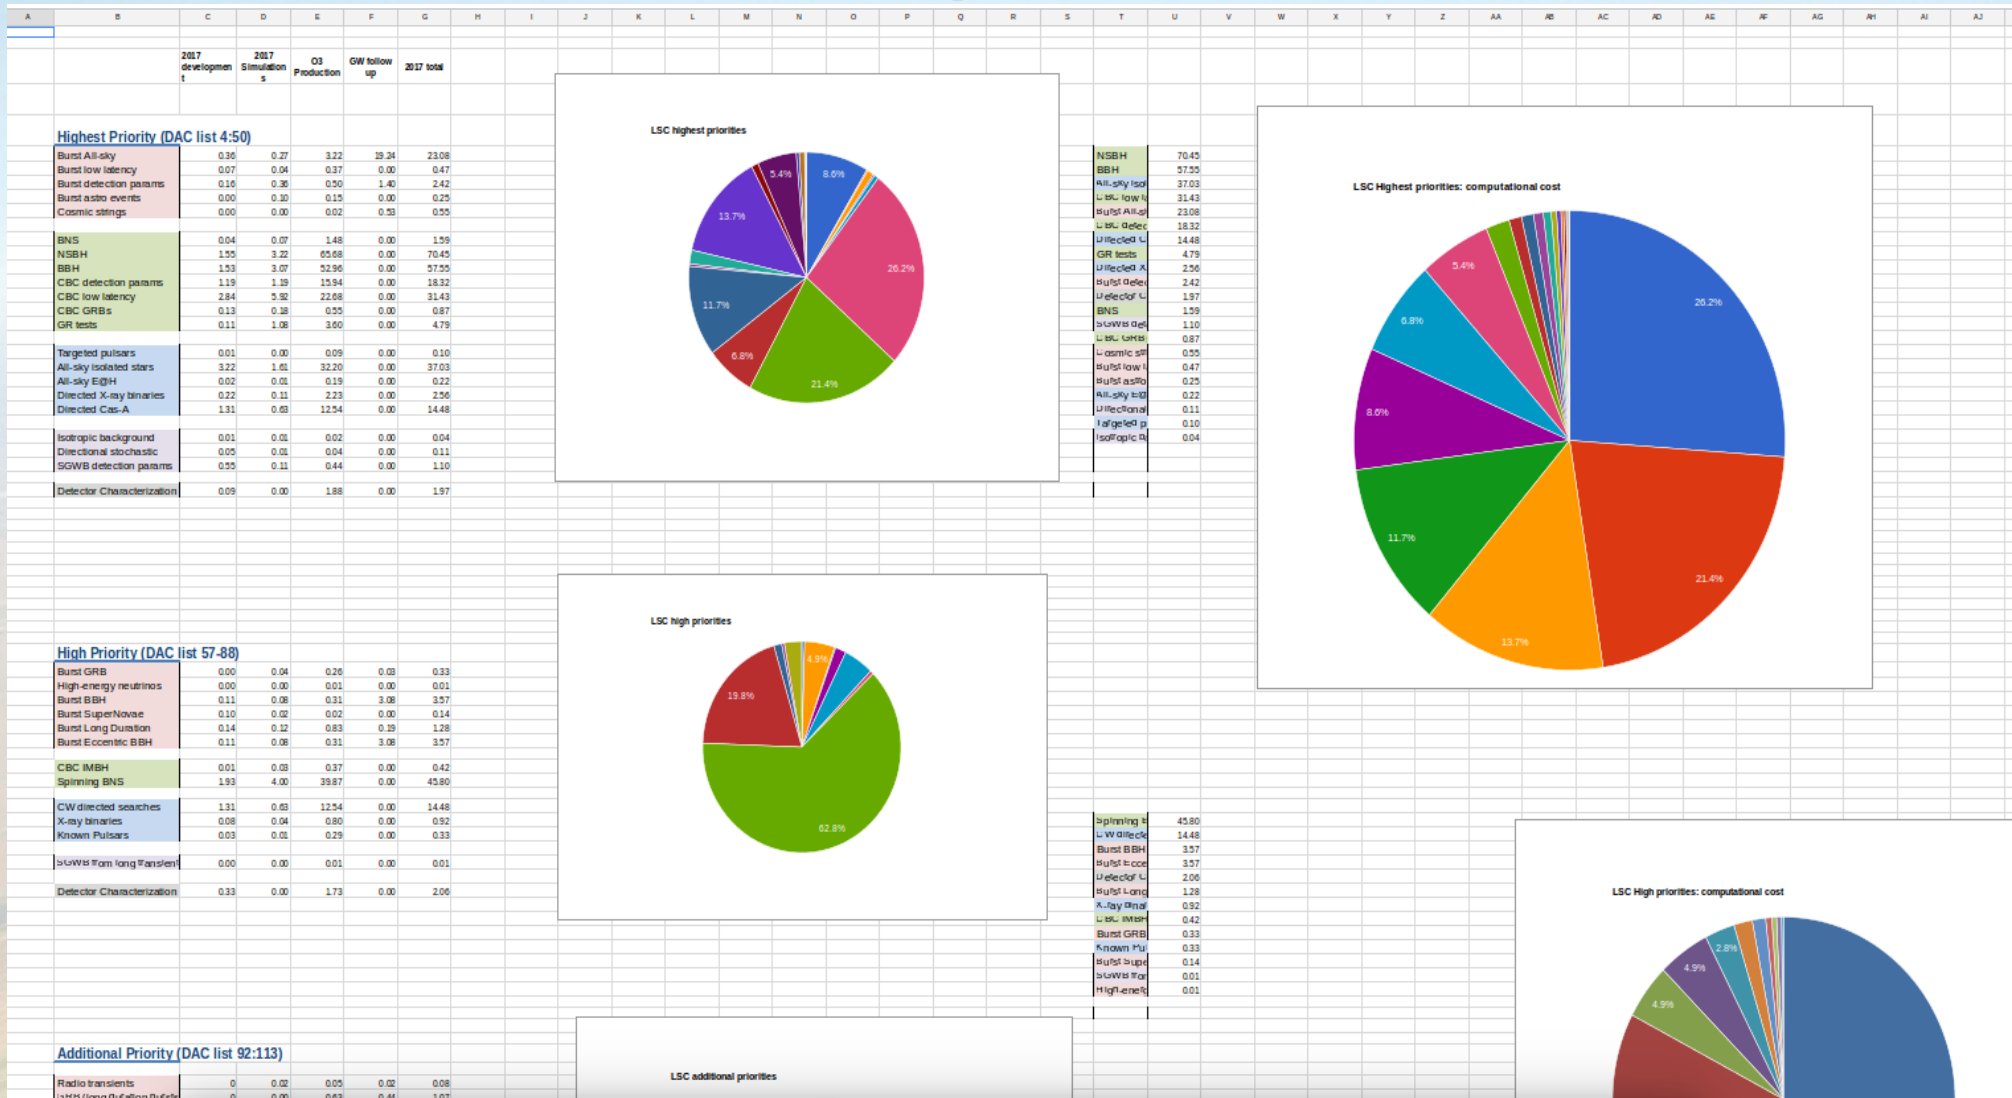
19

An example for pipeline computing requirements collection

2015 may. 25. - INFN CNAF

20

LVCN priorities



Different analysis with different prioritites

Uncertainties of compute and storage estimates

Storage estimates are derived from detector bandwidth and expected duration of observation times known in advance, so there not too much uncertainties here.

As for the compute needs the source of uncertainties are as follows:

- Which pipelines will be ported to GPUs and when
- How successfully the sharing of resources will happen with LIGO collaboration
- Possibility of Condor CLOUD installation
- Available budget for computing in EGO
- Currently Virgo is providing only c.c 8% of the total computing power for the LVC collaboration. This cannot be maintained on the long term, there is a need for massive increase.

The Compute Backend (CB)

The problem

- For several reason (cost efficiency, manpower, future hardwares, etc..) the analysis code has to be generic
- It is always a subject of debate which language to use to program GPUs.
- Double coding for multiple interface is a waste of time and manpower.

The solution:

- **THE COMPUTE BACKEND (CB) IS ADDRESSING THIS PROBLEM BY PROVIDING UNIFIED INTERFACE FOR VARIOUS GPU PROGRAMING LANGUAGES, SUCH AS CUDA AND OPENCL !**
- It levreages the burden of host-side double coding and the very same code can be used to run on CUDA (NVidia) or OpenCL (AMD, Intel, Samsung, etc...) devices...

Compute Backend (CB) features:

- C and C++ API (fortran, python and c# on the way...)
- CUDA and OpenCL backends (ComputeGL, RenderScript considered)
- Single host-side code for multiple backend
- Runs under Linux/Windows/MacOS
- Compatible with CMake, Autoconf, MSVC, etc.
- Academic license is available
- User support around the clock
-

Compute Backend is available on

<http://x-perception.com>

The Compute Backend - the C API

```
#include <stdio.h>
#include <stdlib.h>
#include <cb.h>

int main() {

    // Auxiliary variables
    int err ;
    int i;

    // Sets the log level
    cb_log_level = 5;

    // Get some buffer
    unsigned int num_elements = 1024;
    unsigned int size = num_elements * sizeof(float);

    // ... and also on the host side
    float * h_buffer1 = (float *) malloc(size);
    float * h_buffer2 = (float *) malloc(size);
    float * h_buffer3 = (float *) malloc(size);

    // ... fill up the buffers
    for (i = 0; i < num_elements; i++) {h_buffer1[i] = 4; h_buffer2[i] = 11;}

    // The C API
    // A compute backend
    cb_backend backend;
    cb_program prog;
    cb_kernel kernel1, kernel2, kernel3;
    cb_buffer buffer1, buffer2, buffer3;

    // Get the compute backend
    err = cbGetComputeBackend(&backend);

    // Get a program
    err = cbGetProgram(&backend, "/home/me/testt", &prog);

    // Get the kernel
    err = cbGetKernel(&prog, "test_kernel", &kernel1);
    err = cbGetKernel(&prog, "simple_kernel", &kernel2);
    err = cbGetKernel(&prog, "buffer_kernel", &kernel3);
```

```
err = cbCreateBuffer(&backend, CB_READ_WRITE, size, NULL, &buffer1);
err = cbCreateBuffer(&backend, CB_READ_WRITE, size, NULL, &buffer2);
err = cbCreateBuffer(&backend, CB_READ_WRITE, size, NULL, &buffer3);

    // Send some data to device
    err = cbWriteBuffer(&backend.queues[0], &buffer1, size, h_buffer1, true);
    err = cbWriteBuffer(&backend.queues[0], &buffer2, size, h_buffer2, true);

    // Set the kernel sizes
    cbExtent g_size = cbSetExtent(1,1024);
    cbExtent l_size = cbSetExtent(1, 32);

    // Execute the kernel
    cbParam b1_arg = cbBuffer(&buffer1);
    cbParam b2_arg = cbBuffer(&buffer2);
    cbParam b3_arg = cbBuffer(&buffer3);
    cbParam n_arg = cbInt(100);

    err = cbExecuteKernel(&backend.queues[0], &kernel3, g_size, l_size, 4,
    &b1_arg, &b2_arg, &n_arg, &b3_arg);

    // Read back the result
    err = cbReadBuffer(&backend.queues[0], &buffer3, size, h_buffer3, true);

    // Printing the result
    for (i = 0; i < 10; i++) printf("%f ", h_buffer3[i]);
    printf("\n\n");

    // Releasing stuff
    free(h_buffer1);
    free(h_buffer2);
    free(h_buffer3);

    // Exit
    return err;
}
```


The Compute Backend - the C++ API

```
#include <stdio.h>
#include <stdlib.h>
#include <iostream>
#include <cb.hpp>

int main() {

    // Sets the log level
    cb_log_level = 5;
    int err ;
    int i;

    // Get some buffer on the host side
    unsigned int num_elements = 1024;
    unsigned int size = num_elements * sizeof(float);

    float * h_buffer1 = new float[num_elements];
    float * h_buffer2 = new float[num_elements];
    float * h_buffer3 = new float[num_elements];

    // ... fill in the buffers
    for (i = 0; i < num_elements; i++) {h_buffer1[i] = 4; h_buffer2[i] = 11;}

    // Construction Backend, Program, Kernel and Buffers
    cb::Backend bck;
    cb::Program prg(bck, "/home/me/test");
    cb::Kernel TestKernel(prg, "test_kernel");
    cb::Kernel SimpleKernel(prg, "simple_kernel");
    cb::Kernel BufferKernel(prg, "buffer_kernel");

    // Initializing the buffers
    cb::Buffer b1(bck, CB_READ_WRITE, size, NULL);
    cb::Buffer b2(bck, CB_READ_WRITE, size, NULL);
    cb::Buffer b3(bck, CB_READ_WRITE, size, NULL);

    // Send data to device
    b1.Write(bck.GetQueue(), h_buffer1);
    b2.Write(bck.GetQueue(), h_buffer2);

    // Set the kernel sizes
    cb::Extent g(num_elements);
    cb::Extent l(32);
```

```
// Create kernel arguments
cbParam buff1_arg = cbBuffer(b1);
cbParam buff2_arg = cbBuffer(b2);
cbParam buff3_arg = cbBuffer(b3);
cbParam numarg = cbInt(100);

// Execute the buffer kernel
BufferKernel(bck.GetQueue(), g, l, 4, &buff1_arg, &buff2_arg, &numarg, &buff3_arg);

// Read back the result
b3.Read(bck.GetQueue(), h_buffer3);

// Some output for checking the result
for (int i = 0; i < 10; i++) {
    std::cout << h_buffer1[i] << " " << h_buffer2[i] << " " << h_buffer3[i];
}

// Releasing stuff
delete h_buffer1;
delete h_buffer2;
delete h_buffer3;

// Exiting
exit(0);
}
```

Compile for CUDA:

```
cd build
cmake -DOPENCL_BACKEND=1 ../
make
```

Compile for OpenCL:

```
cd build
cmake -DCUDA_BACKEND=1 ../
make
```

Summary

- Computing model and operation of gravitational wave experiments are well understood and established
- There are multiple problems to be faced and solved due to the usage of loosely coupled resources and distinct administration.
- The community goes under a change, evolution that was experienced in HEP c.c 12 years ago...
- We should use those tools developed and experiences accumulated since then... and avoid the mistakes