Lessons learned in AdV+ phase 1

Report for the EGO Council

Edited and redacted by G. Gemme (Virgo spokesperson 2023-current) under his sole responsibility.

Thanks to G. Losurdo (Virgo spokesperson 2020-2023), F. Sorrentino (Commissioning coordinator 2022-2024), M. Tacca (Commissioning coordinator 2019-2022) and J. van den Brand (Virgo spokesperson 2017-2020) for useful inputs and discussions. Thanks to all other colleagues who commented and helped improve this report.

1. Introduction

This document seeks to retrospectively analyze "*what went wrong*" during phase 1 of Advanced Virgo+ (AdV+): why the initial predictions about detector sensitivity turned out to be wrong, why commissioning took so long, causing Virgo to fail to join at the start of O4, and why some technical problems were not foreseen and whether they could have been identified earlier.

The commissioning of Advanced Virgo Plus has shown the limitations that come from using marginally stable recycling cavities (MSRC) and of the resulting difficulty of operating an interferometer under these conditions. We will not go into too much detail about the technical reasons behind this problem; they are described in detail in other documents (e.g. VIR-0047B-23 and VIR-0026A-24). The current optical configuration of AdV+ certainly plays an important role in the difficulties that the experiment is currently experiencing, but it would be a mistake to reduce the problem to this single cause: multiple factors have contributed, at different levels. Although the impact of marginally stable cavities cannot be underestimated, and the installation of stable cavities in Virgo is a necessary step to ensure the scientific relevance of Virgo in the medium and long term, we made an effort to identify and briefly discuss the technical problems and organizational limitations that we consider most important to understanding the current situation, starting from some considerations on the general context within which the AdV+ project was developed.

2. Context and Background

The term "failure" is often used in connection with phase 1 of AdV+. It is undeniable that the fact that the project sensitivity has not been obtained, that the expected times for the project have expanded well beyond the original plan (at least in the commissioning phase) and that Virgo's decision not to join the initial part of O4 are the manifestation of the failure to achieve fundamental objectives for the project. The results are disappointing also when compared to the ones obtained by LIGO, which is usually a benchmark for Virgo. This comparison with LIGO, while legitimate (these are experiments that have the same scientific goals and are based on largely overlapping technologies), should not be limited to the results obtained by the two experiments, but should include the "boundary conditions" that made it possible to achieve these results. An aspect of these boundary conditions that cannot be ignored is the profound difference in funding, organization and role played within the experiment by the two laboratories that host the observatories: the LIGO laboratory and EGO.

While the LIGO lab's current Maintenance and Operation budget is in the order of $50 million/year, EGO's corresponding budget, despite having recently received a significant increase, still is in the order of €10 million/year[1]. It is clear that an analytical comparison between the two budgets requires an in-depth analysis that takes into account their different organizational structure and the different distribution of responsibilities between the laboratories and scientific collaborations behind the LIGO and Virgo experiments, but it is also clear that the gap in resources is such that it is difficult not to imagine its impact in the long term on the performance of the two detectors.

	Advanced LIGO	Advanced Virgo
# DETECTORS	2+1	1
MAX CBC RANGE (BNS)	200 Mpc	140 Mpc
BUDGET	205[(A)] M$ + 16[(B)] (D/UK/AUS)	21.8[(C)] M€ + 2[(B)] (NL)
FUNDING APPROVED	Apr 2008	Dec 2009
CONSTRUCTION END[(D)]	Jul 2014	May 2016
1st PROJECT REVIEW	2003	2008
MEMBERS	~900	~200
COUNTRIES	17	5
LABS	82	19
R&D INVESTMENTS	~60[(E)] M$	~2[(F)]+1.5[(G)] M€

Table 1. Upgrade investments: Advanced LIGO vs. Advanced Virgo[2]. Snapshot made at the time of the TDR in 2012.

The M&O budget is only part of the story. Looking at the budgets for upgrades and taking into consideration the example of Advanced Virgo (AdV) and Advanced LIGO (see Table 1) it is striking the huge disparity in budget between two projects with similar scientific objectives and based on comparable technologies. The difference of almost an order of magnitude is largely explained by the different accounting used in the United States and Europe; in particular, LIGO's budget includes the cost of personnel dedicated to the project, and moreover it is related to the upgrade of two detectors. However, the fact that personnel dedicated to the upgrade were included in the cost of the project and under the direct responsibility of the laboratory, gives LIGO great flexibility and efficiency in the ability to direct personnel resources and hire new staff according to the needs of the project. This was important in ensuring the success of Advanced LIGO. Finally, the disparity in resources dedicated to investments for R&D, which determine the technical/scientific maturity and reliability of the technologies that make the upgrade possible, speaks for itself and is, unfortunately, a constant gap between our experiments.

---

[1] Of course, the main difference between LIGO and EGO is that LIGO lab handles the operation of two interferometers. This is, of course, a relevant element, but it does not substantially change any of the considerations in the text, in the sense that, for example, the gap in financial resources, even scaled to take into account the presence of two interferometers, remains extremely significant.

[2] A. Includes money for people (“half stuff, half staff”)
B. In kind contribution
C. Only for investments
D. Expected according to the latest planning
E. Personal communication from D Shoemaker. LIGO lab R&D (+2-3 M$/year in other labs)
F. EGO R&D calls 2003 and 2007
G. INFN-CSN2 funding 2005-2010 (data from Fulvio Ricci)

The limited budget related to the upgrade, regarding Advanced Virgo, not only has a "historic" value but has had a direct impact on the AdV+ project. It is well known that the original design of Advanced Virgo was based on an optical configuration with stable recycling cavities and that the decision to change the optical configuration with marginally stable recycling cavities was largely determined by budget limitations.

In addition to the different level of financial resources, in the comparison between LIGO and EGO, it is worth mentioning at least a second aspect, no less relevant than the first: the different "attractiveness" between the two laboratories and the resulting different ability to attract scientists even outside the countries where they are based.

Finally, an element of general context that cannot be ignored is represented by the exponential growth of activities related to the development of third-generation interferometers (Einstein Telescope, ET) in Europe in recent years. The ESFRI schedule, the launch of the Instrument Science Board and of the ET collaboration as early as 2021, the launch of important R&D programs requiring substantial personpower: all these aspects created tensions among a relative small community of experts in gravitational wave detectors instrument science who started to steer their efforts towards ET at the expenses of Virgo. This has made it difficult in some cases to have the availability of key experts in critical phases and, in part, to invest in the formation of new ones[3].

3. The AdV+ Project

Let's start by looking back at the 1-year long study and debate which took place in 2010-11, in the Advanced Virgo (AdV) framework and which eventually led to the decision to run AdV with MSRC. That choice, which anyway allowed to run in O2 and make the breakthrough discovery of GW170817, was based on three main motivations:

- The financial issue, which made it impossible to pursue the long cavity option.
- The risks related to the short vertical cavity solution, which made it hard to pursue it.
- The idea that an advanced TCS could completely tackle the MSRC issues.

Moreover, while a lot of MSRC issues had been anticipated and are summarized in the review report (VIR-0310A-11), others were not: the problems observed during the AdV+ commissioning originating from higher order modes living in the signal recycling (SR) cavity were not known (the simulations done to prepare the implementation of the control system were based on too simple models for an interferometer with marginally stable cavities. As a consequence, this effect was not found in the simulations, probably because of the strongly detuned SR configuration originally planned for AdV).

AdV joined O2 with in power-recycled configuration, the power recycling cavity being marginally stable. The O3 run was carried out with the same configuration and increased sensitivity. The proposal of the AdV+ upgrade was born surfing the wave of the success of O2 and O3. In the "AdV+ Design Report" (VIR-0596A-19) there is no discussion on the risk of continuing with MSRC (MSRC are mentioned just once in 400 pages, in the TCS chapter). The idea apparently was that no problems were expected with phase 1. There are some risks with phase 2 and to mitigate them

[3] It has also to be taken in consideration that the "external interference" can have an impact on the mid-term plan: the realization of stable cavities with the target of joining part of O5 requires focus, extensive commitment and dedication by the whole collaboration.

one must further develop the TCS. A thorough and full simulation study of the dual recycled configuration was not done, also due to the lack of suitable tools to carry them out. The risk with MSRC was also overlooked by the International Review Committee (there is no mention in the report to the Council of July 2019).

We can say that the AdV+ project was born with an "original sin":

- the lack of detailed, realistic, simulations,
- the fact that, due to this and the good results obtained in O2 and O3, no clear showstoppers emerged and, finally,
- the fact that no one wanted to face again a long and divisive debate like the one that happened in 2010-11,

led to the underestimation of the impact of MSRCs in AdV+. On top of this,

- the lack of a "Plan B", included in the commissioning plan, to be put in place in case of setbacks, to ensure that Virgo could join O4 on time and with decent sensitivity was not available from the start of the project.

4. Lessons learned during commissioning

The commissioning after AdV+ phase 1 upgrades took longer than expected; Virgo missed the start of the O4 science run, and the current sensitivity is between the levels of O3a and O3b.

The interferometer is now operated with a circulating optical power at the level of O2 (18 W of input power), and the best BNS range (around 50 Mpc) is obtained with a trade-off between detector bandwidth and gain by reducing the level of signal recycling (SR) with a slight misalignment of the SR mirror.

In such conditions, the sensitivity around 100 Hz is basically limited by quantum noise, coating thermal noise (to a lesser extent), and an unmodeled "mystery" noise with weak frequency dependence. We know that a mystery noise component was already limiting the Virgo sensitivity during O3, although now there is no decisive and unambiguous indication that these two noises originate from the same source.

Here we try to give some considerations about the lessons learned from O4 commissioning, and about the main reasons for the delay and lower performance with respect to the original plan.

4.1. Detector upgrades

Impact of hardware. The preparation of critical devices during detector upgrade can affect the commissioning duration; in general, important installations should be possibly anticipated, as the impact of hardware interventions on the schedule during commissioning can be much higher than during the installation or pre-commissioning phases. Despite a successful installation/integration phase, completed on schedule during the pandemic, there are a few examples of issues that had an impact on the O4 commissioning schedule:

- Some monitoring hardware for thermal compensation, planned for AdV+ phase 2, was later realized to be needed already in phase 1 commissioning; the installation and setup required several weeks of re-tuning and lock recovery.

- The lack of spares for some critical hardware generated a few weeks of delay during O4 commissioning: this happened e.g. with the fibered laser amplifier for O4 that was found to be defective, with the backup option already used in O3 damaged and lacking a working spare.
- Some issues came from the aging of hardware components. An example is DSP electronics: due to the excellent performance of the superattenuators in O3, an upgrade of their driving electronics for O4 was considered unnecessary despite a request from the subsystem. The boards are naturally aging, and our stock of spares is getting low: frequent faults in DSP boards for suspension control occurred, sometimes producing elusive noises. Another example is the input laser control electronics and opto-electronics: the laser frequency control loop and photodiode readout, based on old analog electronics, are rather critical and require some delicate retuning from time to time. Such kind of issues caused several weeks of delay in the commissioning.
- Following the outcome of an internal review, it was chosen to run O4 with the same set of test masses used in O3, having damaged monolithic suspensions, and not to replace them with spares having lower mechanical and optical losses; in the end we had more than two months of commissioning breaks to fix issues with mechanical losses on test masses. Replacing one of the test masses during commissioning revealed a substantial improvement in optical and mechanical losses, even if no change in sensitivity was observed.
- Some devices, not ready before the start of O4 commissioning, were installed later, e.g. power stabilization electronics, residual amplitude modulation servo, point absorbers mitigation system; the overall impact on commissioning schedule was a delay of more than 2 months. For example, breaking the vacuum in the test mass vacuum chamber to install the point absorber mitigation system caused the failure of one monolithic suspension; the repair of the suspension and the recovery of the system required more than one month.
- It was known that an actuator for the compensation of a cold optical defect in core optics, i.e. the ring heater for power recycling (PR) mirror curvature, had been mounted incorrectly, and was therefore ineffective, since O2. The late installation of a backup actuator generated a delay of about one year in O4 commissioning, as no reproducible stable lock was possible before the installation on the PR of the central heating radius of curvature correction (CHRoCC) in July 2022. More generally, the presence of large cold defects requires the use of large corrections with thermal actuators, which can in turn generate important optical aberrations. Whenever possible, removing the major cold defects (currently the PR mirror curvature and the cold lens in the north input compensation plate) may be a relevant improvement.
- The stability of some subsystems proved to be vulnerable, and sensitive to external factors; this applies i.e. to Power Stabilized Laser/Injection (PSL/INJ), where slight tuning of parameters (e.g. changes on input laser optical power), or just the presence of people in laser lab, can trigger periods of instability or excess noise; and to the Auxiliary Laser System, where the sensitivity to environmental noises can severely limit the lock duty cycle in periods of bad weather and/or during lock recovery. Issues with such critical subsystems carried off a few weeks of commissioning time overall.

Impact of marginally stable recycling cavities. The commissioning of AdV+ has shown how difficult it is to operate the interferometer with marginally stable recycling cavities. While some of these difficulties had been anticipated by simulations at the time of the design of Advanced Virgo, they were considered workable and for this reason Advanced Virgo was built with marginally stable recycling cavities. The mitigation strategy identified at the time of the design of Advanced Virgo was based on an advanced thermal compensation. Indeed it was possible to operate Advanced Virgo in the power recycled configuration and participate in O2 and O3. No showstoppers were found at that time. However, the interferometer was operated with 25 W input power instead of the 125 W foreseen at the time of the Advanced Virgo design.

The commissioning of AdV+ started in January 2021. The implementation of the signal recycling cavity made the interferometer operation considerably more complex. After the first lock at 25 W of input power, a long, stable lock with 40 W input power was not possible. For this reason the power was decreased to 33 W. The first two hour lock with 33 W of input power were achieved at the end of 2021.

After several studies and the addition of an actuator able to change the radius of curvature of the recycling mirrors (CHRoCC), the turning point was reached in the fall of 2022 when it was found that the signal recycling cavity was resonating high order modes and so worsening the interferometer contrast defect. This has several consequences:

- Optical tuning is difficult and time consuming because the poor selectivity to optical modes limits the accuracy of optical characterization, and because of the resulting extreme complexity of optical simulations does not allow accurate modeling.
- Bistability of interferometer lock is observed for a large set of optical parameters, which adds further complexity and slowness to optical tunings.
- A larger fraction of the laser power reaches the detection system thus making its operation more difficult.
- The interferometer is more sensitive to laser noises.
- There are large losses in the signal recycling cavity that affect the performance of the squeezing system.

This becomes worse when the power is increased even at power levels well below the Advanced Virgo design.

It is clear that operating the interferometer with MSRC represents a challenge, and that this difficulty will be worse at increasing optical power, as the contribution of high order optical modes and the requirements on thermal compensation scale with the power; this sets a possible indirect sensitivity limitation from MSRC, because it limits the power one can inject into the interferometer, which is however still not exactly quantified. On the other hand we did not find any decisive experimental evidence that the “mystery” current noise currently limiting the sensitivity might be directly due to MSRC.

4.2. Commissioning strategy

Lack of focus on noise hunting. Commissioning after detector upgrades should have two main goals: stable operation of the interferometer in the new configuration, and mitigation of technical noises; the latter should ideally take a larger fraction of commissioning time, as noise hunting is the most delicate task. Once the first goal is achieved, the interferometer should run in DC readout for most of the time, to identify and address the most offending noise sources. Control optimization in RF readout is not particularly useful, as it delays the critical activity in the final configuration of the interferometer (DC readout) where the impact on sensitivity can be studied. This happened in part during O4 commissioning, where noise hunting took less than 30 weeks (i.e. 20% of the available time), while we spent up to 14 weeks on controls optimization in RF readout.

Optical defects can have a large impact on the interferometer performance, both in terms of control stability and noise couplings, and must be taken into due account.

- Cold optical defects must be precisely compensated; as explained above, the lack of proper compensation for a cold defect on PR mirror curvature prevented the reproducible stable lock of the interferometer for a large fraction of the O4 commissioning (nearly one year).
- Thermal effects depend on both the circulating optical power and on defects in core optics (e.g. point absorbers). To disentangle thermal effects from the impact of cold defects and of optical alignments, the interferometer should be first commissioned at low circulating power, and then at progressively higher levels of optical power. The choice to run at O4 design optical power at an early stage of the commissioning phase, in combination with the missing correction of the power recycling mirror cold defect, produced an important delay in the O4 commissioning and did not allow to quantify the impact of optical power on interferometer performance. It was not possible to achieve a stable lock at 40 W input power, and stable interferometer control has been achieved with an input power up to 33 W. The current input power (February 2024), 18 W, has been chosen in order to optimize the thermal compensation system in a simpler configuration. No further increase is foreseen before the start of O4b due to the limited time available for commissioning.

4.3. Commissioning organization and management

- The absence of a written, widely discussed and agreed commissioning plan has been a weakness all along and especially in the initial phase of the commissioning.
- Commissioning planning should try to make best use of human resources and machine time. This is often the result of a difficult balance to avoid requiring excessive flexibility from available personpower. Optimal use of machine time would require to have as many commissioning shifts as possible; on the other hand, expert commissioners are precious human resources and should not be overloaded, which may result in loss of motivation and efficiency. The combination of longer duration of O4 commissioning, the small effective size of the commissioning crew on site (also due to the difficulty of integrating new people into the team) and the pressure to maximally exploit machine time with 13 shifts per week eventually produced negative effects on the team, including lack of motivation, focus and efficiency.
- Commissioning planning should be done with due involvement of committed experts from the involved subsystems. This was not the case during the first half of O4 commissioning and produced a lack of focus and motivation from many of the most involved commissioners.
- The separation between the responsibilities of the Project Leader and of the Commissioning Coordinator in the distinct phases of the project was not clear at the beginning, creating tensions. Confusion of roles and interference of project management in commissioning planning generated misunderstandings, lack of motivation and lack of efficiency in the commissioning work. This happened during the first two years of O4

commissioning, with the AdV+ management directly micro-managing individual commissioning shifts in the control room.

- A management team including the figures of “System Managers”, intermediate between the Project Leader and the Subsystem managers, was created at the beginning of the project. Although this structure may have favored coordination and communication within the same System, communication and the definition of interfaces between Subsystems belonging to different Systems has not always been optimal. In addition, the definition of responsibilities between System and Subsystem managers has not always been clear, sometimes causing misunderstandings and tensions.
- At an advanced stage of commissioning, a task force of former commissioning coordinators and few experts was set-up by the project management, with the approval of the VSC. This action proved to be only partially successful, both because the availability of some task force members to interact with continuity with the commissioning team has been limited, and because of resistance from commissioners to accept outside "interference."
- Personpower is critical on some subsystems where a single expert is currently available. In such conditions a continuous on-call service is not possible, and in the occurrence of emergency events (e.g. power outages) this generated substantial delays in system recovery. In addition, this generated a large pressure on on-site EGO experts, which are in fact committed to cover personpower vacancies from laboratories without any formal mandate.

5. Quality assurance/quality control

Insufficient effort in QA/QC and system engineering was already highlighted by the STAC in their November 2022 report: ”Continuous effort has been made to enhance system engineering activities. Quality Assurance and Quality Control is part of it and is addressed through reviews. The STAC nevertheless feels that this is perhaps not sufficient. It recommends for example to Investigate the possibility to implement for EGO a non-conformity / incidents / … tracking and management software which would be a tool to support QA/QC.”

While a substantial effort was put in place by the System Engineer and others in order to put in place document controls, reviews, inspection and test planning that proved to be effective during the course of the Project (see e.g. VIR-0540C-22), other good practices such as the use of workmanship standards for AdV+ equipment and facilities have essentially failed. “Workmanship Standards”, imply a collaboration level, large and multidisciplinary effort that is currently largely lacking. For this reason, the risk is that some of the System Engineering threads become idle because the involved people (reviewers, subsystem managers, etc.) are not responsive enough (mostly because “too busy” in “doing” the project). Unless the overall project management (or more in general the expert staff engaged in the upgrade) is strongly reinforced we are probably at the limit of how much QA&QC we can “inject”.

6. Summary and possible actions

The answer to the question "*What went wrong?*" in Phase 1 of AdV+ is complex, but some causes can be identified and some possible actions can be suggested :

- Tight budgets and minimized upgrade investments (the current crisis is, in part, a result of choices made in the past).
- Competition with Einstein Telescope, and lack of global coordination at the level of the funding agencies. To mitigate this tendency there is absolute need for coordination at a level above EGO/Virgo and ET, at the level of funding agencies, for a strategic plan that involves both Virgo_nEXT and ET, laying out the sharing of funding and personpower.
- The relatively modest resources, financial and personnel, allocated to EGO for the operation of a detector of increasing complexity.
- The poor “attractiveness” of EGO for senior and junior scientists. Making EGO more attractive for scientists to work side-by-side at the site for a longer time should become a priority. Virgo experts need to be on site to engage in discussions with the on-site experts from EGO and with collaboration and project management; everybody would profit from daily and continuous exchange with senior experts.
- Critical personpower situation in some laboratories of the Virgo collaboration with crucial hardware responsibilities.
- Small effective size of the commissioning crew on site.
- Lack of a commissioning plan and lack of a “Plan B” to be put in place in case of setbacks from the beginning of the project; unclear and suboptimal commissioning strategy; unclear separation of responsibilities and decision-making process.
- Lack of accurate simulations of a realistic configuration of the detector to prepare the implementation of the control system, also caused by the lack of adequate simulation tools.
- The underestimation of the impact of MSRCs in AdV+ and the resulting impact on commissioning.
- Several hardware issues that required ad-hoc and time consuming interventions.
- Insufficient effort in QA/QC and system engineering. The lack of a collaboration level culture of workmanship standards.

The installation of stable cavities in Virgo is a necessary step to ensure the scientific relevance of Virgo in the medium and long term, as well as its role as a pathfinder for the Einstein Telescope. However, it should not be considered the only action to be taken if the problems highlighted during phase 1 of the AdV+ project are to be avoided. A wide-ranging action and an assumption of responsibility by all the actors in the field is needed: the laboratories of the Virgo collaboration, the EGO laboratory and the funding agencies.