

PAPER

Virgo detector characterization and data quality: results from the O3 run


























To cite this article: F Acernese *et al* 2023 *Class. Quantum Grav.* **40** 185006

View the [article online](#) for updates and enhancements.

You may also like

- [The ATLAS Fast Tracker system](#)
The ATLAS collaboration, G. Aad, B. Abbott *et al.*




Virgo detector characterization and data quality: results from the O3 run

F Acernese^{1,2}, M Agathos³, A Ain⁴, S Albanesi^{5,6},
 A Allocca^{7,2} , A Amato⁸ , T Andrade⁹, N Andres¹⁰ ,
 M Andrés-Carcasona¹¹ , T Andrić¹² , S Ansoldi^{13,14},
 S Antier^{15,16} , T Apostolatos¹⁷, E Z Appavuravther^{18,19},
 M Arène²⁰, N Arnaud^{21,22,*} , M Assiduo^{23,24},
 S Assis de Souza Melo²², P Astone²⁵ , F Aubin²⁴ ,
 S Babak²⁰ , F Badaracco²⁶ , M K M Bader²⁷,
 S Bagnasco⁶ , J Baird²⁰, T Baka²⁸, G Ballardín²²,
 G Baltus²⁹ , B Banerjee¹² , C Barbieri^{30,31,32}, P Barneo⁹ ,
 F Barone^{33,2} , M Barsuglia²⁰ , D Barta³⁴ , A Basti^{35,4},
 M Bawaj^{18,36} , M Bazzan^{37,38}, F Beirnaert³⁹ , M Beijger⁴⁰ ,
 I Belahcene²¹, V Benedetto⁴¹, M Berbel⁴² , S Bernuzzi³ ,
 D Bersanetti⁴³ , A Bertolini²⁷, U Bhardwaj^{16,27} ,
 A Bianchi^{27,44}, S Bini^{45,46}, M Bischì^{23,24}, M Bitossi^{22,4},
 M-A Bizouard¹⁵ , F Bobba^{47,48}, M Boër¹⁵, G Bogaert¹⁵,
 M Boldrini^{49,25}, L D Bonavena³⁷, F Bondu⁵⁰, R Bonnand¹⁰ ,
 B A Boom²⁷, V Boschi⁴ , V Boudart²⁹ , Y Bouffanais^{37,38},
 A Bozzi²², C Bradaschia⁴, M Branchesi^{12,51} , M Breschi³ ,
 T Briant⁵² , A Brillat¹⁵, J Brooks²², G Bruno²⁶, F Buccì²⁴,
 T Bulik⁵³, H J Bulten²⁷, D Buskulic¹⁰, C Buy⁵⁴ ,
 G S Cabourn Davies⁵⁵ , G Cabras^{13,14} , R Cabrita²⁶ ,
 G Cagnoli⁸ , E Calloni^{7,2}, M Canepa^{56,43}, S Canevarolo²⁸,
 M Cannavacciuolo⁴⁷, E Capocasa²⁰ , G Carapella^{47,48},
 F Carbognani²², M Carpinelli^{57,58,22}, G Carullo^{35,4} ,
 J Casanueva Diaz²², C Casentini^{59,60}, S Caudill^{27,28},
 F Cavalier²¹ , R Cavalieri²² , G Cella⁴ ,
 P Cerdá-Durán⁶¹, E Cesarini⁶⁰ , W Chaibi¹⁵,
 P Chaniál²² , E Chassande-Mottin²⁰ , S Chaty²⁰ ,
 F Chiadini^{62,48} , G Chiarini³⁸, R Chierici⁶³, A Chincarini⁴³ ,
 M L Chiofalo^{35,4}, A Chiummo²² , S Choudhary⁶⁴ ,
 N Christensen¹⁵ , G Ciani^{37,38} , P Cielag⁴⁰,
 M Cieślak⁴⁰ , M Cifaldi^{59,60}, R Cioffi^{65,38} , F Cipriano¹⁵,
 S Clesse⁶⁶, F Cleva¹⁵, E Coccia^{12,51}, E Codazzo¹² ,
 P-F Cohadon⁵² , D E Cohen²¹ , A Colombo^{30,31} ,
 M Colpi^{30,31}, L Conti³⁸ , I Cordero-Carrión⁶⁷ , S Corezzi^{36,18},

* Author to whom any correspondence should be addressed.

D Corre²¹, S Cortese²² , J-P Coulon¹⁵, M Croquette⁵² ,
 J R Cudell²⁹ , E Cuoco^{22,68,4}, M Curyło⁵³, P Dabadie⁸,
 T Dal Canton²¹ , S Dall’Osso¹² , G Dálya³⁹ ,
 B D’Angelo^{56,43} , S Danilishin^{69,27} , S D’Antonio⁶⁰,
 V Dattilo²², M Davier²¹, D Davis⁷⁰ , J Degallaix⁷¹ ,
 M De Laurentis^{7,2}, S Deléglise⁵² , F De Lillo²⁶ ,
 D Dell’Aquila⁵⁷ , W Del Pozzo^{35,4}, F De Matteis^{59,60},
 A Depasse²⁶ , R De Pietri^{72,73} , R De Rosa^{7,2} ,
 C De Rossi²², R De Simone⁶², L Di Fiore², C Di Giorgio^{47,48} ,
 F Di Giovanni⁶¹ , M Di Giovanni¹², T Di Girolamo^{7,2} ,
 A Di Lieto^{35,4} , A Di Michele³⁶ , S Di Pace^{49,25} ,
 I Di Palma^{49,25} , F Di Renzo^{35,4} , L D’Onofrio^{7,2} ,
 M Drago^{49,25} , J-G Ducoin²¹, U Dupletsa¹², O Durante^{47,48},
 D D’Urso^{57,58} , P-A Duverne²¹, M Eisenmann¹⁰,
 L Errico^{7,2}, D Estevez⁷⁴ , F Fabrizi^{23,24} , F Faedi²⁴,
 V Fafone^{59,60,12} , S Farinon⁴³, G Favaro³⁷ , M Fays²⁹ ,
 E Fenyvesi^{34,75} , I Ferrante^{35,4} , F Fidocarò^{35,4} ,
 P Figura⁵³ , A Fiori^{4,35} , I Fiori²² , R Fittipaldi^{76,48},
 V Fiumara^{77,48}, R Flaminio^{10,78}, J A Font^{61,79} , S Frasca^{49,25},
 F Frasconi⁴ , A Freise^{27,44} , O Freitas⁸⁰, G G Fronzé⁶ ,
 B U Gadre^{81,28} , R Gamba³, B Garaventa^{43,56} , F Garufi^{7,2} ,
 G Gemme⁴³ , A Gennai⁴ , Archisman Ghosh³⁹ ,
 B Giacomazzo^{30,31,32} , L Giacoppo^{49,25}, P Giri^{4,35} ,
 F Gissi⁴¹, S Gkaitatzis^{4,35} , B Goncharov¹² , M Gosselin²²,
 R Gouaty¹⁰, A Grado^{82,2} , M Granata⁷¹ , V Granata⁴⁷,
 G Greco¹⁸, G Grignani^{36,18}, A Grimaldi^{45,46} , S J Grimm^{12,51},
 P Gruning²¹, D Guerra⁶¹ , G M Guidi^{23,24} , G Guixé⁹,
 Y Guo²⁷, P Gupta^{27,28}, L Haegel²⁰ , O Halim¹⁴ ,
 O Hannuksela^{28,27}, T Harder¹⁵, K Haris^{27,28}, J Harms^{12,51} ,
 B Haskell⁴⁰, A Heidmann⁵² , H Heitmann¹⁵ , P Hello²¹,
 G Hemming²² , E Hennes²⁷ , S Hild^{69,27}, D Hofman⁷¹,
 V Hui¹⁰ , B Idzkowski⁵³ , A Iess^{59,60}, P Iosif⁸³ ,
 T Jacqmin⁵² , P-E Jacquet⁵² , S P Jadhav⁶⁴,
 J Janquart^{28,27}, K Janssens^{84,15} , P Jaranowski⁸⁵ ,
 V Juste⁷⁴, C Kalaghatgi^{28,27,86}, C Karathanasis¹¹ ,
 S Katsanevas^{22,†} , F Kéfélian¹⁵, N Khetan^{12,51},
 G Koekoek^{27,69}, S Koley¹² , M Kolstein¹¹ , A Królak^{87,88} ,
 P Kuijjer²⁷ , P Lagabbe¹⁰, D Laghi⁵⁴ , M Lalleman⁸⁴,
 A Lamberts^{15,89}, I La Rosa¹⁰, A Lartaux-Vollard²¹,
 C Lazzaro^{37,38}, P Leaci^{49,25} , A Lemaître⁹⁰,
 M Lenti^{24,91} , E Leonova¹⁶, N Leroy²¹ , N Letendre¹⁰,
 K Leyde²⁰, F Linde^{86,27}, L London¹⁶, A Longo⁹² ,

† Deceased.

M Lopez Portilla²⁸, M Lorenzini^{59,60} , V Lorient⁹³,
 G Losurdo⁴ , D Lumaca^{59,60} , A Macquet¹⁵, C Magazzù⁴ ,
 M Magnozzi^{43,56} , E Majorana^{49,25}, I Maksimovic⁹³,
 N Man¹⁵, V Mangano^{49,25} , M Mantovani²² , M Mapelli^{37,38} ,
 F Marchesoni^{19,18,94}, D Marín Pina⁹ , F Marion¹⁰,
 A Marquina⁶⁷, S Marsat²⁰ , F Martelli^{23,24}, M Martinez¹¹,
 V Martinez⁸, A Masserot¹⁰, S Mastrogiovanni²⁰ , Q Meijer²⁸,
 A Menendez-Vazquez¹¹, L Mereni⁷¹, M Merzougui¹⁵,
 A Miani^{45,46} , C Michel⁷¹ , L Milano^{7,†}, A Miller²⁶,
 B Miller^{16,27}, E Milotti^{95,14}, Y Minenkov⁶⁰, LI M Mir¹¹,
 M Miravet-Tenés⁶¹ , M Montani^{23,24}, F Morawski⁴⁰,
 B Mours⁷⁴ , C M Mow-Lowry^{27,44} , S Mozzon⁵⁵ ,
 F Muciaccia^{49,25}, Suvodip Mukherjee¹⁶ , R Musenich^{43,56} ,
 A Nagar^{6,96}, V Napolano²², I Nardecchia^{59,60} , H Narola²⁸,
 L Naticchioni²⁵, J Neilson^{41,48}, C Nguyen²⁰ , S Nissanke^{16,27},
 E Nitoglia⁶³ , F Nocera²², G Oganessian^{12,51}, C Olivetto²²,
 G Pagano^{35,4}, G Pagliaroli^{12,51}, C Palomba²⁵ ,
 P T H Pang^{27,28}, F Pannarale^{49,25} , F Paoletti⁴ , A Paoli²²,
 A Paolone^{25,97}, G Pappas⁸³, D Pascucci^{27,39} ,
 A Pasqualetti²², R Passaquieti^{35,4} , D Passuello⁴,
 B Patricelli^{22,4} , R Pedurand⁴⁸, M Pegoraro³⁸, A Perego^{45,46},
 A Pereira⁸, C Périgois¹⁰, A Perreca^{45,46} , S Perriès⁶³,
 D Pesios⁸³, K S Phukon^{27,86} , O J Piccinni²⁵ , M Pichot¹⁵ ,
 M Piendibene^{35,4}, F Piergiovanni^{23,24}, L Pierini^{49,25} ,
 V Pierro^{41,48} , G Pillant²², M Pillas²¹, F Pilo⁴, L Pinard⁷¹,
 I M Pinto^{41,48,98}, M Pinto²², K Piotrkowski²⁶, A Placidi^{18,36} ,
 E Placidi^{49,25}, W Plastino^{99,92} , R Poggiani^{35,4} , E Polini¹⁰ ,
 E K Porter²⁰, R Poulton²² , M Pracchia¹⁰, T Pradier⁷⁴,
 M Principe^{41,98,48}, G A Prodi^{100,46} , P Proposito^{59,60},
 A Puecher^{27,28}, M Punturo¹⁸ , F Puosi^{4,35}, P Puppo²⁵,
 G Raaijmakers^{16,27}, N Radulesco¹⁵, P Rapagnani^{49,25},
 M Razzano^{35,4} , T Regimbau¹⁰, L Rei⁴³ , P Rettegno^{5,6} ,
 B Revenu²⁰ , A Reza²⁷, F Ricci^{49,25}, G Riemenschneider^{5,6},
 S Rinaldi^{35,4} , F Robinet²¹, A Rocchi⁶⁰ , L Rolland¹⁰ ,
 M Romanelli⁵⁰, R Romano^{1,2}, A Romero¹¹ ,
 S Ronchini^{12,51} , L Rosa^{2,7}, D Rosińska⁵³, S Roy²⁸,
 D Rozza^{57,58} , P Ruggi²², J Sadiq¹⁰¹ , O S Salafia^{32,31,30} ,
 L Salconi²², F Salemi^{45,46} , A Samajdar³¹ ,
 N Sanchis-Gual¹⁰² , A Sanuy⁹ , B Sassolas⁷¹, S Sayah⁷¹,
 S Schmidt²⁸, M Seglar-Arroyo¹⁰ , D Sentenac²²,
 V Sequino^{7,2}, Y Setyawati²⁸ , A Sharma^{12,51},
 N S Shcheblanov⁹⁰ , M Sieniawska²⁶, L Silenzi^{18,19} ,

† Deceased.

N Singh⁵³ , **A Singha**^{69,27} , **V Sipala**^{57,58},
J Soldateschi^{91,103,24} , **K Soni**⁶⁴ , **V Sordini**⁶³,
F Sorrentino⁴³, **N Sorrentino**^{35,4} , **R Soulard**¹⁵,
V Spagnuolo^{69,27}, **M Spera**^{37,38} , **P Spinicelli**²²,
C Stachie¹⁵, **D A Steer**²⁰ , **J Steinlechner**^{69,27},
S Steinlechner^{69,27} , **N Stergioulas**⁸³, **G Stratta**^{104,25} ,
M Suchenek⁴⁰, **A Sur**⁴⁰ , **B L Swinkels**²⁷ , **P Szcwzyk**⁵³,
M Tacca²⁷, **A J Tanasijczuk**²⁶, **E N Tapia San Martín**²⁷ ,
C Taranto⁵⁹, **A E Tolley**⁵⁵ , **M Tonelli**^{35,4}, **A Torres-Forné**⁶¹ ,
I Tosta e Melo⁵⁸ , **A Trapananti**^{19,18} , **F Travasso**^{18,19} ,
M Trevor¹⁰⁵ , **M C Tringali**²² , **L Troiano**^{106,48}, **A Trovato**²⁰ ,
L Trozzo², **K W Tsang**^{27,107,28}, **K Turbang**^{108,84} , **M Turconi**¹⁵,
A Utina^{69,27} , **M Valentini**^{45,46} , **N van Bakel**²⁷,
M van Beuzekom²⁷ , **M van Dael**^{27,109},
J F J van den Brand^{69,44,27} , **C Van Den Broeck**^{28,27},
H van Haeveermaet⁸⁴ , **J V van Heijningen**²⁶ ,
N van Remortel⁸⁴ , **M Vardaro**^{86,27}, **M Vasúth**³⁴ ,
G Vedovato³⁸, **D Verkindt**¹⁰ , **P Verma**⁸⁸, **F Vetrano**²³,
A Vicere^{23,24} , **V Villa-Ortega**¹⁰¹ , **J-Y Vinet**¹⁵,
A Virtuoso^{95,14}, **H Vocca**^{36,18}, **R C Walet**²⁷, **M Was**¹⁰ ,
A R Williamson⁵⁵ , **J L Willis**⁷⁰ , **A Zadrożny**⁸⁸,
T Zelenova²² and **J-P Zendri**³⁸

¹ Dipartimento di Farmacia, Università di Salerno, I-84084 Fisciano, Salerno, Italy

² INFN, Sezione di Napoli, Complesso Universitario di Monte S. Angelo, I-80126 Napoli, Italy

³ Theoretisch-Physikalisches Institut, Friedrich-Schiller-Universität Jena, D-07743 Jena, Germany

⁴ INFN, Sezione di Pisa, I-56127 Pisa, Italy

⁵ Dipartimento di Fisica, Università degli Studi di Torino, I-10125 Torino, Italy

⁶ INFN Sezione di Torino, I-10125 Torino, Italy

⁷ Università di Napoli “Federico II”, Complesso Universitario di Monte S. Angelo, I-80126 Napoli, Italy

⁸ Université de Lyon, Université Claude Bernard Lyon 1, CNRS, Institut Lumière Matière, F-69622 Villeurbanne, France

⁹ Institut de Ciències del Cosmos (ICCUB), Universitat de Barcelona, C/ Martí i Franquès 1, Barcelona 08028, Spain

¹⁰ Univ. Savoie Mont Blanc, CNRS, Laboratoire d’Annecy de Physique des Particules—IN2P3, F-74000 Annecy, France

¹¹ Institut de Física d’Altes Energies (IFAE), Barcelona Institute of Science and Technology, and ICREA, E-08193 Barcelona, Spain

¹² Gran Sasso Science Institute (GSSI), I-67100 L’Aquila, Italy

¹³ Dipartimento di Scienze Matematiche, Informatiche e Fisiche, Università di Udine, I-33100 Udine, Italy

¹⁴ INFN, Sezione di Trieste, I-34127 Trieste, Italy

¹⁵ Artemis, Université Côte d’Azur, Observatoire de la Côte d’Azur, CNRS, F-06304 Nice, France

- ¹⁶ GRAPPA, Anton Pannekoek Institute for Astronomy and Institute for High-Energy Physics, University of Amsterdam, Science Park 904, 1098 XH Amsterdam, The Netherlands
- ¹⁷ Department of Physics, National and Kapodistrian University of Athens, School of Science Building, 2nd floor, Panepistimiopolis, 15771 Ilissia, Greece
- ¹⁸ INFN, Sezione di Perugia, I-06123 Perugia, Italy
- ¹⁹ Università di Camerino, I-62032 Camerino, Italy
- ²⁰ Université de Paris, CNRS, Astroparticule et Cosmologie, F-75006 Paris, France
- ²¹ Université Paris-Saclay, CNRS/IN2P3, IJCLab, 91405 Orsay, France
- ²² European Gravitational Observatory (EGO), I-56021 Cascina Pisa, Italy
- ²³ Università degli Studi di Urbino “Carlo Bo”, I-61029 Urbino, Italy
- ²⁴ INFN, Sezione di Firenze, I-50019 Sesto Fiorentino Firenze, Italy
- ²⁵ INFN, Sezione di Roma, I-00185 Roma, Italy
- ²⁶ Université catholique de Louvain, B-1348 Louvain-la-Neuve, Belgium
- ²⁷ Nikhef, Science Park 105, 1098 XG Amsterdam, The Netherlands
- ²⁸ Institute for Gravitational and Subatomic Physics (GRASP), Utrecht University, Princetonplein 1, 3584 CC Utrecht, The Netherlands
- ²⁹ Université de Liège, B-4000 Liège, Belgium
- ³⁰ Università degli Studi di Milano-Bicocca, I-20126 Milano, Italy
- ³¹ INFN, Sezione di Milano-Bicocca, I-20126 Milano, Italy
- ³² INAF, Osservatorio Astronomico di Brera sede di Merate, I-23807 Merate Lecco, Italy
- ³³ Dipartimento di Medicina, Chirurgia e Odontoiatria “Scuola Medica Salernitana”, Università di Salerno, I-84081 Baronissi Salerno, Italy
- ³⁴ Wigner RCP, RMKI, H-1121 Budapest, Konkoly Thege Miklós út 29-33, Hungary
- ³⁵ Università di Pisa, I-56127 Pisa, Italy
- ³⁶ Università di Perugia, I-06123 Perugia, Italy
- ³⁷ Università di Padova, Dipartimento di Fisica e Astronomia, I-35131 Padova, Italy
- ³⁸ INFN, Sezione di Padova, I-35131 Padova, Italy
- ³⁹ Universiteit Gent, B-9000 Gent, Belgium
- ⁴⁰ Nicolaus Copernicus Astronomical Center, Polish Academy of Sciences, 00-716 Warsaw, Poland
- ⁴¹ Dipartimento di Ingegneria, Università del Sannio, I-82100 Benevento, Italy
- ⁴² Departamento de Matemáticas, Universitat Autònoma de Barcelona, Edificio C Facultad de Ciencias 08193 Bellaterra (Barcelona), Spain
- ⁴³ INFN, Sezione di Genova, I-16146 Genova, Italy
- ⁴⁴ Vrije Universiteit Amsterdam, 1081 HV Amsterdam, The Netherlands
- ⁴⁵ Università di Trento, Dipartimento di Fisica, I-38123 Povo Trento, Italy
- ⁴⁶ INFN, Trento Institute for Fundamental Physics and Applications, I-38123 Povo Trento, Italy
- ⁴⁷ Dipartimento di Fisica “E.R. Caianiello”, Università di Salerno, I-84084 Fisciano, Salerno, Italy
- ⁴⁸ INFN, Sezione di Napoli, Gruppo Collegato di Salerno, Complesso Universitario di Monte S. Angelo, I-80126 Napoli, Italy
- ⁴⁹ Università di Roma “La Sapienza”, I-00185 Roma, Italy
- ⁵⁰ Univ Rennes, CNRS, Institut FOTON—UMR6082, F-35000 Rennes, France
- ⁵¹ INFN, Laboratori Nazionali del Gran Sasso, I-67100 Assergi, Italy
- ⁵² Laboratoire Kastler Brossel, Sorbonne Université, CNRS, ENS-Université PSL, Collège de France, F-75005 Paris, France
- ⁵³ Astronomical Observatory Warsaw University, 00-478 Warsaw, Poland
- ⁵⁴ L2IT, Laboratoire des 2 Infinis—Toulouse, Université de Toulouse, CNRS/IN2P3, UPS, F-31062 Toulouse Cedex 9, France
- ⁵⁵ University of Portsmouth, Portsmouth PO1 3FX, United Kingdom

- ⁵⁶ Dipartimento di Fisica, Università degli Studi di Genova, I-16146 Genova, Italy
- ⁵⁷ Università degli Studi di Sassari, I-07100 Sassari, Italy
- ⁵⁸ INFN, Laboratori Nazionali del Sud, I-95125 Catania, Italy
- ⁵⁹ Università di Roma Tor Vergata, I-00133 Roma, Italy
- ⁶⁰ INFN, Sezione di Roma Tor Vergata, I-00133 Roma, Italy
- ⁶¹ Departament de Astronomia y Astrofísica, Universitat de València, E-46100 Burjassot, València, Spain
- ⁶² Dipartimento di Ingegneria Industriale (DIIN), Università di Salerno, I-84084 Fisciano, Salerno, Italy
- ⁶³ Université Lyon, Université Claude Bernard Lyon 1, CNRS, IP2I Lyon/IN2P3, UMR 5822, F-69622 Villeurbanne, France
- ⁶⁴ Inter-University Centre for Astronomy and Astrophysics, Post Bag 4, Ganeshkhind, Pune 411 007, India
- ⁶⁵ INAF, Osservatorio Astronomico di Padova, I-35122 Padova, Italy
- ⁶⁶ Université libre de Bruxelles, Avenue Franklin Roosevelt 50, 1050 Bruxelles, Belgium
- ⁶⁷ Departament de Matemàtiques, Universitat de València, E-46100 Burjassot, València, Spain
- ⁶⁸ Scuola Normale Superiore, Piazza dei Cavalieri, 7 - 56126 Pisa, Italy
- ⁶⁹ Maastricht University, P.O. Box 616, 6200 MD Maastricht, The Netherlands
- ⁷⁰ LIGO Laboratory, California Institute of Technology, Pasadena, CA 91125, United States of America
- ⁷¹ Université Lyon, Université Claude Bernard Lyon 1, CNRS, Laboratoire des Matériaux Avancés (LMA), IP2I Lyon/IN2P3, UMR 5822, F-69622 Villeurbanne, France
- ⁷² Dipartimento di Scienze Matematiche, Fisiche e Informatiche, Università di Parma, I-43124 Parma, Italy
- ⁷³ INFN, Sezione di Milano Bicocca, Gruppo Collegato di Parma, I-43124 Parma, Italy
- ⁷⁴ Université de Strasbourg, CNRS, IPHC UMR 7178, F-67000 Strasbourg, France
- ⁷⁵ Institute for Nuclear Research, Bem t'er 18/c, H-4026 Debrecen, Hungary
- ⁷⁶ CNR-SPIN, c/o Università di Salerno, I-84084 Fisciano, Salerno, Italy
- ⁷⁷ Scuola di Ingegneria, Università della Basilicata, I-85100 Potenza, Italy
- ⁷⁸ Gravitational Wave Science Project, National Astronomical Observatory of Japan (NAOJ), Mitaka City, Tokyo 181-8588, Japan
- ⁷⁹ Observatori Astronòmic, Universitat de València, E-46980 Paterna, València, Spain
- ⁸⁰ Centro de Física das Universidades do Minho e do Porto, Universidade do Minho, Campus de Gualtar, PT-4710-057 Braga, Portugal
- ⁸¹ Max Planck Institute for Gravitational Physics (Albert Einstein Institute), D-14476 Potsdam, Germany
- ⁸² INAF, Osservatorio Astronomico di Capodimonte, I-80131 Napoli, Italy
- ⁸³ Department of Physics, Aristotle University of Thessaloniki, University Campus, 54124 Thessaloniki, Greece
- ⁸⁴ Universiteit Antwerpen, Prinsstraat 13, 2000 Antwerpen, Belgium
- ⁸⁵ University of Białystok, 15-424 Białystok, Poland
- ⁸⁶ Institute for High-Energy Physics, University of Amsterdam, Science Park 904, 1098 XH Amsterdam, The Netherlands
- ⁸⁷ Institute of Mathematics, Polish Academy of Sciences, 00656 Warsaw, Poland
- ⁸⁸ National Center for Nuclear Research, 05-400 Świerk-Otwock, Poland
- ⁸⁹ Laboratoire Lagrange, Université Côte d'Azur, Observatoire Côte d'Azur, CNRS, F-06304 Nice, France
- ⁹⁰ NAVIER, École des Ponts, Univ Gustave Eiffel, CNRS, Marne-la-Vallée, France

- ⁹¹ Università di Firenze, Sesto Fiorentino I-50019, Italy
⁹² INFN, Sezione di Roma Tre, I-00146 Roma, Italy
⁹³ ESPCI, CNRS, F-75005 Paris, France
⁹⁴ School of Physics Science and Engineering, Tongji University, Shanghai 200092, People's Republic of China
⁹⁵ Dipartimento di Fisica, Università di Trieste, I-34127 Trieste, Italy
⁹⁶ Institut des Hautes Etudes Scientifiques, F-91440 Bures-sur-Yvette, France
⁹⁷ Consiglio Nazionale delle Ricerche—Istituto dei Sistemi Complessi, Piazzale Aldo Moro 5, I-00185 Roma, Italy
⁹⁸ Museo Storico della Fisica e Centro Studi e Ricerche “Enrico Fermi”, I-00184 Roma, Italy
⁹⁹ Dipartimento di Matematica e Fisica, Università degli Studi Roma Tre, I-00146 Roma, Italy
¹⁰⁰ Università di Trento, Dipartimento di Matematica, I-38123 Povo Trento, Italy
¹⁰¹ Instituto Galego de Física de Altas Enerxías, Universidade de Santiago de Compostela, 15782 Santiago de Compostela, Spain
¹⁰² Departamento de Matemática da Universidade de Aveiro and Centre for Research and Development in Mathematics and Applications, Campus de Santiago, 3810-183 Aveiro, Portugal
¹⁰³ INAF, Osservatorio Astrofisico di Arcetri, Largo E. Fermi 5, I-50125 Firenze, Italy
¹⁰⁴ Istituto di Astrofisica e Planetologia Spaziali di Roma, Via del Fosso del Cavaliere, 100, 00133 Roma RM, Italy
¹⁰⁵ University of Maryland, College Park, MD 20742, United States of America
¹⁰⁶ Dipartimento di Scienze Aziendali—Management and Innovation Systems (DISA-MIS), Università di Salerno, I-84084 Fisciano Salerno, Italy
¹⁰⁷ Van Swinderen Institute for Particle Physics and Gravity, University of Groningen, Nijenborgh 4, 9747 AG Groningen, The Netherlands
¹⁰⁸ Vrije Universiteit Brussel, Pleinlaan 2, 1050 Brussel, Belgium
¹⁰⁹ Eindhoven University of Technology, Postbus 513, 5600 MB Eindhoven, The Netherlands

E-mail: nicolas.arnaud@ijclab.in2p3.fr

Received 31 October 2022; revised 25 March 2023

Accepted for publication 26 May 2023

Published 14 August 2023



CrossMark

Abstract

The Advanced Virgo detector has contributed with its data to the rapid growth of the number of detected GW signals in the past few years, alongside the two Advanced LIGO instruments. First during the last month of the Observation Run 2 (O2) in August 2017 (with, most notably, the compact binary mergers GW170814 and GW170817), and then during the full Observation Run 3 (O3): an 11 months data taking period, between April 2019 and March 2020, that led to the addition of 79 events to the catalog of transient GW sources maintained by LIGO, Virgo and now KAGRA. These discoveries and the manifold exploitation of the detected waveforms benefit from an accurate characterization of the quality of the data, such as continuous study and monitoring of the detector noise sources. These activities, collectively named *detector characterization and data quality* or *DetChar*, span the whole workflow of the Virgo data, from the instrument front-end hardware to the final analyses. They are described in

detail in the following article, with a focus on the results achieved by the Virgo DetChar group during the O3 run. Concurrently, a companion article describes the tools that have been used by the Virgo DetChar group to perform this work.

Keywords: gravitational waves, Advanced Virgo detector, detector characterization, data quality, Observing run 3

(Some figures may appear in colour only in the online journal)

Contents

1. Introduction	10
2. The Advanced Virgo detector	11
2.1. The path to the O3 run	12
2.2. The O3 configuration	13
2.3. Virgo data and DetChar products	15
3. The O3 run	17
3.1. Organization	17
3.1.1. Data taking	17
3.1.2. Detector steering	18
3.1.3. DetChar organization and tools	19
3.1.4. On call duty service and RRT meetings	20
3.2. Performance	20
3.2.1. Noise budget	20
3.2.2. Virgo O3 duty cycle	23
3.2.3. The Virgo O3 dataset	26
4. Real-time data quality	27
4.1. The Virgo O3 online data quality framework	28
4.1.1. State vector	29
4.1.2. Online CAT1 vetoes	29
4.1.3. SegOnline	29
4.2. Veto streams	30
5. Public alerts	32
5.1. Performance of the Virgo O3 DQR framework	32
5.2. O3 public alerts	34
5.2.1. Public alerts retracted because of an issue with Virgo data	34
5.2.2. Virgo contribution to O3 public alerts	35
6. Global data quality studies	36
6.1. Glitches and pipeline triggers	36
6.1.1. Glitch rates during the O3 run	36
6.1.2. Offline searches	37
6.2. Channel safety: channel (in)sensitivity to gravitational waves	40
6.3. Spectral noise	41
6.3.1. Combs	42
6.3.2. Wandering line around 83 Hz–84 Hz	43
6.3.3. Spectral noise bump around 55 Hz	44
6.3.4. Spectral noise around the 50 Hz power line frequency	44
6.4. Offline data quality	45
6.4.1. Offline studies and checks	45

6.4.2. Event validation	47
7. Preparation of the O4 run	47
Data availability statement	48
Acknowledgments	48
References	54

List of abbreviations

AdV	Advanced Virgo
ASD	amplitude spectral density
BH	black hole
BNS	binary neutron star
BRMS	band-limited RMS
BruCo	brute-force coherence tool
BS	beam splitter
CARM	common (i.e. average) length of the two arm cavities
CEB	central building
CW	continuous gravitational waves
DAQ	data acquisition system
DARM	difference of the two arm cavity lengths
DMS	Detector Monitoring System
DOF	degree of freedom
DQ	data quality
DQR	Data Quality Report
DQSEGDB	Data Quality SEGment Database
EOM	electro-optical modulator
FFT	fast Fourier transform
GraceDB	GRAvitational-wave Candidate Event DataBase
GW	gravitational wave
GWOSC	Gravitational Wave Open Science Center
IMC	input mode-cleaner
LVA1ert	LIGO-Virgo Alert System
MICH	length difference between the Virgo Michelson interferometer short arms
NE	north end
NEB	north-end building
NI	north input
NS	neutron star
OMC	output mode-cleaner
PR	power recycling
PRCL	power recycling cavity length
RRT	rapid-response team
SGWB	stochastic gravitational-wave background
SNEB	suspended north-end bench
SNR	signal-to-noise ratio
SR	signal recycling
SSFS	second-stage frequency stabilization system
SWEB	suspended west-end bench
UPV	Use-Percentage Veto
VIM	Virgo Interferometer Monitor
WE	west end
WEB	west-end building
WI	west input

1. Introduction

A century after being predicted by Albert Einstein in the framework of general relativity, GWs have been detected by a global network of ground-based interferometric detectors [1]. The LIGO [2] and Virgo [3] collaborations, now joined by the KAGRA [4] collaboration, have observed in the past seven years dozens of GW signals coming from merging compact binary systems. Compact binaries composed of two BHs, two NSs, or both kinds of compact object have now all been observed. GW150914 [1], the first GW signal ever detected (at that time by the two Advanced LIGO detectors only) was a binary BH merger. Two years later, shortly after the AdV detector had started operating, the LIGO-Virgo 3-interferometer network detected the signal GW170817 [5], emitted by the fusion of two NSs and associated with counterparts in the entire electromagnetic spectrum, leading to the birth of multi-messenger astronomy with GWs. More recently, LIGO, Virgo and KAGRA (in short ‘LVK’) have announced the first detections of NS-BH mergers in data taken in January 2020 [6].

All these events add up in a GW Transient Catalog whose first four versions—GWTC-1 [7], GWTC-2 [8], GWTC-2.1 [9] and GWTC-3 [10]—have been successively released. Such catalogs allow scientists to analyze all detections globally: to probe the populations of compact stars, estimate the merger rate of binary systems, test general relativity in the strong-field regime, and perform searches for counterparts using archival data from other observatories. The reconstructed GW strain data—the so-called $h(t)$ streams—are regularly released in chunks of several months on the GWOSC website [11].

Producing these results requires a thorough characterization of the data quality, a large part of which involves studying and monitoring the noise of GW detectors. This activity, some aspects of which are often referred to as *detector characterization* or *DetChar*, is an expertise which has been constructed over many years, starting more than two decades ago, first on simulated [12] or prototype [13] data, then with the initial detectors [14, 15]. Analysis methods and tools have been developed and implemented to characterize the Virgo data both with low latency¹¹⁰ and offline. The various analyses cover a physics-driven chain that starts from the raw data recorded by the instrument and extends all the way to the final set of GW events and the related analyses. The results of DetChar studies are used to improve the detector performances during the commissioning periods and to maximize the sensitivity to GW signals during an *Observation Run*, when good-quality data are recorded. They are also used to define the final Virgo dataset—used by the LVK analyses and later published on the GWOSC website—, and to vet all GW candidates, found either in low-latency or offline.

This article reports the work of the Virgo DetChar group over the past few years. It will mainly focus on the activities carried out for the preparation of the third LIGO-Virgo observation run O3, from April 2019 to March 2020 as well as on the final results achieved after the run. The experience accumulated in view of the future runs of the LVK network will also be presented. These achievements stem from the developments made before and during the O2 run (for Virgo: 25 days of data taking in August 2017) and those will also be described here when appropriate.

¹¹⁰ Throughout this article, the word *latency* is used to indicate the period between when an information becomes available (for instance some data quality assessment) and the time when the corresponding data have been acquired by the Virgo detector.




Figure 1. Flowchart of the Virgo DetChar tools and monitors described in [16]. The corresponding data quality (in short ‘DQ’) products are used to monitor the detector and as inputs to GW searches.

Concurrently, a companion article [16] describes in detail all the tools the Virgo DetChar group has been relying on, in order to obtain the results presented here.

This article is organized as follows. Section 2 provides an overview of the AdV detector configuration during the O3 run, preceded by a short summary of the path that led to this data taking period. The same section also introduces notions and concepts that will be extensively used in the rest of the article, and defines a few related abbreviations. Then, section 3 summarizes the O3 run from a Virgo perspective: how the data taking was organized, what the performance of the detector and the final O3 dataset were. Section 4 presents the Virgo online data quality framework built for the O3 run. Section 5 deals with the software developed to vet signal candidates to be released as public alerts to the astronomical community. Section 6 presents the main DetChar analyses done on the O3 dataset to study noise transients, their impact on GW searches, the noise spectrum, and the final validation of events. Finally, section 7 provides some information about the ongoing preparation of the future O4 run, that is currently scheduled to begin in Spring 2023. A list of the main abbreviations used throughout the article is provided as well, for reference.

Finally, figure 1 (a simplified version of the flowchart used in [16] to show an overall view of the DetChar tools) summarizes the Virgo data flow and describes how the Virgo DetChar products are combined to provide inputs to the GW searches.

2. The Advanced Virgo detector

This section focuses on the Virgo detector during the O3 run. First, we briefly review the main steps of the AdV project up to the beginning of O3. In particular, we emphasize its participation to the last four weeks of the O2 run in August 2017 that were rich of discoveries. Then, we summarize the activities during the 1.5 year-long shutdown between O2 and O3 that allowed the Virgo Collaboration to improve the instrument significantly. Finally, we describe the detector configuration during O3 and present the main features of the data it has collected.

2.1. The path to the O3 run

Virgo [17] is an interferometric detector of GWs located at the European Gravitational Observatory (EGO) in Cascina, Italy. The AdV project [3] allowed to upgrade the original instrument to a second-generation detector, similarly to what LIGO has done with its two interferometers [2], located in Hanford (WA, USA) and Livingston (LA, USA). The funding of AdV was approved in December 2009 by CNRS (France) and INFN (Italy), with an in-kind contribution from Nikhef (The Netherlands). The decommissioning of the first-generation Virgo detector started in Fall 2011, after the completion of the science run VSR4 [18], pursued together with the GEO600 detector [19] (the Advanced LIGO upgrade project had already started). The installation of the Advanced Virgo equipment started mid-2012 and was completed in 2016. The upgraded interferometer was robustly controlled in March 2017 and the next few months were dedicated to commissioning activities: noise hunting and sensitivity improvement. At the end of July, the detector was very stable and had a sensitivity corresponding to a BNS range¹¹¹ of ~ 30 Mpc, that is more than a factor two above the performance of the Virgo+ detector during the VSR4 run.

Therefore, AdV started taking data on 1 August 2017, joining the second Observing Run O2, which had started on 30 November 2016 for the two LIGO interferometers [20]. On 14 August 2017, the AdV detector made its first detection of a GW. That event, labeled GW170814 [21], was also recorded by the two LIGO interferometers. It was the first ever triple detection of a binary black hole coalescence, allowing an unprecedented accuracy in the localization of the source in the sky. A few days later, on August 17, the three interferometers jointly detected, for the first time, a GW signal emitted by the coalescence of two neutron stars [5]. This event, known as GW170817, was accompanied by the almost simultaneous detection of a gamma-ray burst by the Fermi Gamma-ray and Integral space telescopes [22]. The accuracy in the localization of the GW source (approx. 30deg^2) allowed to identify the optical counterpart in the galaxy NGC4993 [23]. The O2 run ended on 25 August 2017.

The LIGO-Virgo shutdown between O2 and the third Observation Run O3, lasted 19 months. On the Virgo side, it was divided into four periods:

- A post-O2 commissioning phase, until early December 2017. The goal was twofold: to make a series of measurements on the O2 detector configuration that would have been too invasive during the run, and to perform some tests to try to further improve the instrument.
- Hardware upgrades until mid-March 2018. Four main projects were pursued:

* **The re-installation of the mirror suspensions and various vacuum upgrades.**

The steel wires, with which the AdV arm cavity mirrors were suspended for the O2 run, were replaced with quartz fibers in order to reduce the friction at the mirror-wire contact points—a source of thermal noise. Fused silica ‘monolithic’¹¹² suspensions, successfully tested in the Virgo+ configuration [24, 25], were foreseen in the AdV Technical Design Report [26]. Yet, multiple breakages of fused silica fibers when installed in vacuum were observed during Fall 2016, forcing the recourse of steel wires to preserve the participation of Virgo to the O2 run. The fiber breaking issue was eventually demonstrated to be caused

¹¹¹ The BNS range is the average distance up to which the merger of a BNS system can be detected. The average is taken over the source location in the sky and the BNS system orientation, while a detection is defined as a signal-to-noise ratio of 8 or above.

¹¹² The wires are welded to the mirrors and both are composed of the same fused silica.

by a spurious dust contamination generated by some vacuum pumps [27]. Therefore, the Virgo vacuum system was improved in order to avoid dust contamination, while the suspension fibers were shielded to prevent them from being hit by dust particles set in motion by ‘air flows’ when acting on the vacuum system.

* **A higher laser power.**

The power of the laser injected into the interferometer was increased, reducing the photon shot noise that is limiting the high-frequency sensitivity: 10 W (19 W) were injected in Virgo during the O2 run (at the beginning of the O3 run).

* **The installation of a squeezed light source.**

This allows to further reduce the shot noise limit at high frequencies by modifying the quantum properties of the light coming out of the interferometer [28].

* **The test installation of an array of seismic sensors.**

An in-depth characterization of the seismic noise field at the test mass locations was performed in order to prepare for the subtraction of the Newtonian noise contribution that may limit the low-frequency sensitivity in the future [29].

- A commissioning period, until Fall 2018, to improve the sensitivity and the duty cycle of the detector.
- Finally, the transition phase to the O3 run, that officially started on 1 April 2019 at 15:00 UTC.

2.2. The O3 configuration

The AdV detector [26, 30] has been designed to achieve a sensitivity about one order of magnitude better compared to the initial Virgo detector, corresponding to an increase in the detection rate by about three orders of magnitude. The AdV design choices were made on the basis of the outcome of the different research and development activities carried out within the GW community and the experience gained with initial Virgo, while also taking into account budget and schedule constraints.

The simplified optical schematic of AdV during the O2 and O3 runs is shown in figure 2. In the following, we briefly outline the different parts of the detector layout and define the main abbreviations that are labeled on the schematic or used later in the article. Further information about the O3 configuration and control system of the Virgo detector can be found in [31].

The Virgo power-stabilized infrared laser beam (wavelength: $1.064 \mu\text{m}$) is filtered at the interferometer input by a 144 m triangular cavity called the IMC. The two flat mirrors of the IMC are located on the first suspended injection bench, that also hosts various optics for beam matching¹¹³. Then, the beam goes through the partially reflective PR mirror before being split into two perpendicular beams at the BS mirror. The two 3 km-long arms hosting Fabry–Perot cavities are called ‘North’ and ‘West’ as they are roughly oriented along these geographical directions. The cavity mirrors closest (furthest away) from the BS are called ‘input’ (‘end’) mirrors. So, following these conventions, the test masses (the four mirrors forming the two 3-km long Fabry–Perot cavities) are labeled NI, NE, WI and WE, respectively. Both arms end with a suspended terminal bench—called SNEB or SWEB—hosting a photodiode (B7 or B8) receiving the cavity transmitted beam. More generally, most optical components and sensors are located on suspended benches (not displayed on the schematic) to reduce the impact of the residual seismic motion. After propagation and storage in the kilometric cavities, the arm beams recombine on the BS and the beam resulting from this interference goes to the interferometer

¹¹³ Making the laser beam wavefront matching the one resonating in the arm cavities.

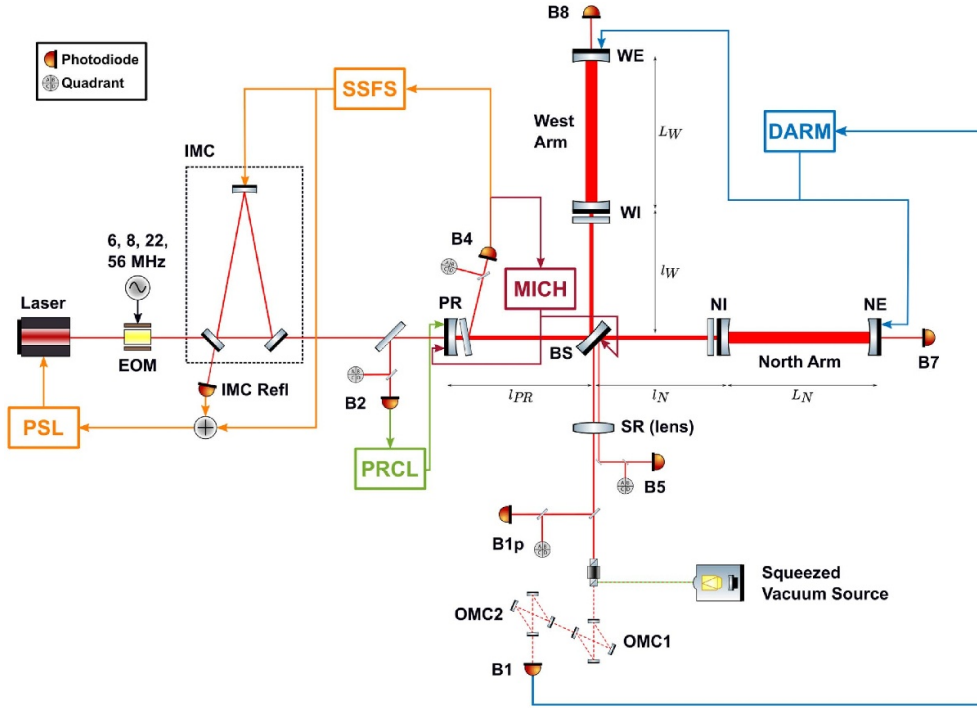


Figure 2. Schematics of the AdV configuration during the O3 run (not to scale), showing optics, photodiodes and quadrant photodiodes, such as the main components of the global feedback system used to steer the detector. The suspended optical benches introduced in the text are not represented here. Reproduced from [31]. CC BY 4.0.

output port. As indicated in figure 2, the location of the foreseen SR mirror was occupied by the first lens of the detection system during the O3 run (and during O2 as well).¹¹⁴ The beam from the frequency-independent squeezed light source [28] enters the detector between the SR lens and the interferometer output. Finally, prior to being detected on the B1 photodiode located on the suspended detection bench 2 (SDB2), the output port beam is filtered in sequence by two output mode-cleaner (OMC) cavities, OMC1 and OMC2, located on the SDB1.

A complex active feedback system, made of several automated control feedback loops, is necessary to bring and maintain the detector at its global working point. In particular, it aims at controlling the four main longitudinal DOFs of the AdV detector which, in its O2–O3 configuration (see figure 2 for the definition of the different lengths used below), are:

- The MICH, $l_N - l_W$, sets the destructive interference (‘dark fringe’) optimal condition.
- The PRCL, $l_{PR} + (l_N + l_W)/2$, must be resonant.
- The lengths of the kilometric Fabry–Perot cavities, L_N and L_W , must be resonant as well, or rather their average and difference that are more physical:
 - * The CARM, $(L_N + L_W)/2$, used as a length etalon by the SSFS to stabilize further the frequency of the input laser.
 - * The DARM, $L_N - L_W$, the quantity sensitive to a passing GW.

¹¹⁴ The installation of that additional mirror only took place during the shutdown period that followed the end of O3.

This global control relies on radio-frequency sidebands for the carrier beam that are generated by the EOM located in between the laser source and the IMC on figure 2. The 6, 8 and 56 MHz sidebands¹¹⁵ are used to control the interferometer, while the 22 MHz one is used to control the injection system.

2.3. Virgo data and DetChar products

The GW strain data stream $h(t)$ reconstructed at the Virgo detector is dominated by noise with, up-to-now, rare and weak GW signals. The noise contributions can be roughly classified into two main categories:

- Fundamental noises, that are inherent to the instrument and represent the ultimate limit of its sensitivity. Their combined contribution is overall stationary and Gaussian [14], two properties that are thoroughly verified as part of the DetChar activities, and in particular in correspondence of candidate GW events [16].
- Various noise artifacts, whose origins are manifold (hardware components of the detector, feedback control loops, interaction with the external environment, etc) and that represent potential issues, not only because they may impact the running of the instrument but also—and above all—because they show up in the background of searches for GW signals, thus limiting their sensitivity. Noise transients, called *glitches*, can either look like real signals or overlap in time with one, either impairing its detection or confusing the inference of its source parameters. These glitches are monitored and studied with time-frequency representations that are used to classify their numerous signatures into families and separate them from real GW events. In addition, long-lasting noise excesses, also called *spectral noises*, are also seen around particular frequencies (power main frequency and its harmonics, suspension resonating modes, etc): the narrow ones, (nearly) monochromatic, are called *lines* and the wider ones *bumps*. Both can manifest themselves in several ‘flavours’. For instance, lines can exist individually, but sometimes appear as *combs*, that is families of lines separated by a constant frequency interval. They are typically due to processes with a strict time periodicity, like electronic clock signals. Bumps may have some specific structure, depending on the source. Both lines and bumps can exhibit structures symmetric around their main frequency, called *sidebands*, that are due to non-linear interactions among different disturbances. Moreover, spectral noises can be persistent across a full run, or only be present in a portion of it. Both the glitch rate in a particular frequency band and the properties (amplitude, peak frequency and bandwidth) of spectral noises can vary in time, to reflect changes occurring at the level of the detector or its environment.

To allow investigating these variations, hundreds of *auxiliary channels* are acquired by the Virgo DAQ, providing both a detailed status of the detector control systems and a complete

¹¹⁵ The exact frequencies of the different sidebands are provided for instance in [26]; they have been chosen to make these sidebands resonant / anti-resonant / non-resonant in the different optical cavities, depending on their usage by the Virgo global control system.




Figure 3. Generic workflow of transient noise studies in Virgo, using the full set of DetChar and monitoring tools described in [16].

monitoring of the local environment [32, 33]. Integer GPS ranges used to flag data with common properties (data quality level, particular detector condition, etc) are called *segments* in the following.

The Virgo GW strain data $h(t)$ and the many associated auxiliary channels are analyzed by a wide set of DetChar tools described in detail in [16]. As an example, figure 3 describes the joint application of various DetChar analysis and monitoring tools to the study of transient noise. This flowchart focuses specifically on how these tools are used and complement each other to investigate transient noise. Omicron is the main tool to identify glitches in all the relevant DAQ channels. Those glitch triggers are stored on disk and mined by the UPV tool, to look for coincidences between them, allowing to assign confidently a terrestrial origin to a fraction of these triggers. VetoPerf monitors the performance of these tools to find an optimum balance between the fraction of glitches flagged as non-astrophysical and the amount of data removed by these associations. Other tools like BRMSMon look for patterns in the data that are known to be due to noise. All these inputs are used to trigger further noise investigations and are gathered to allow a global assessment of the quality of the data. The hypothesis of stationary (and Gaussian) data are also tested around all GW candidates as they are basic assumptions for the algorithms searching for GWs in the data, whether they are run in real-time or offline. Parallel to this dataflow, dedicated monitoring tools like the DMS and the VIM continuously provide information about the status of all the Virgo components, from the hardware blocks to the online software processes.

3. The O3 run

The joint LIGO-Virgo Observing Run 3 O3, has been divided into two consecutive sub-data-taking periods, separated by a 1 month commissioning break in October 2019:

- O3a: from 1 April 2019 at 15:00 UTC (GPS: 12 381 66 018), to 1 October 2019 at 15:00 UTC (GPS: 1253 977 218).
- O3b: from 1 November 2019 at 15:00 UTC (GPS: 12 566 55 618), to 27 March 2020 at 17:00 UTC (GPS: 1269 363 618).

All three detectors have participated to the whole run. The O3b end date has been moved forward by more than a month, due to the worldwide Covid-19 pandemic.

This section presents the LIGO-Virgo O3 run, seen from a Virgo perspective. First, we describe the main activities into which the data taking was divided, before summarizing how the detector was steered from the EGO control room. Then, we focus on actions taken to maximize the amount of data collected and to ensure their good quality. In particular, we highlight the main DetChar activities during O3, explaining how they fit and complement each other, following the flow of data from the detector to the final analyses. Key to achieve this level of performance and to maintain it over almost a year, were the 24/7 on-call duty service and the RRT: both will be briefly described as well.

Then, we review the performance of the Virgo detector during O3, mainly from the point of view of the duty cycle. A high duty cycle requires not only a stable and robust detector against external disturbances (see [33] for a comprehensive study of that topic) but also a quick and reliable procedure to bring the instrument to its working point, starting from an uncontrolled global state. The main statistics of the Virgo O3 global control acquisition are thus provided, before studying the actual duty cycle. We also present the evolution of the AdV detector sensitivity, from the O2 run to the end of O3.

This section ends with a brief overview of the final Virgo O3 dataset, describing how it was constructed offline, building upon the preliminary dataset established by the live monitoring and data quality checks.

3.1. Organization

3.1.1. Data taking. While data acquisition was the highest priority during the O3 run, a limited fraction of the time had to be dedicated to other activities. The two main recurring ones were:

- The maintenance periods, held every Tuesday morning, staggered with respect to the similar times in LIGO, in order to maximize the two-detector network coverage. Maintenance, limited to about 4 h per week, was used to look after the detector components, to perform various cleaning actions, and to host noisy activities incompatible with data taking—for instance the refilling of liquid nitrogen tanks located nearby the CEB, NEB and WEB, delivered by heavy trucks.
- The calibration shifts, held almost every week on Wednesday afternoons or evenings. These campaigns allowed to check the accuracy of the reconstruction of the $h(t)$ stream [34], to monitor its stability over time and to test new, complimentary calibration methods, like the use of a Newtonian calibration system [35] in addition to the usual photon calibrators [36].




Figure 4. Metatron nodes hierarchy used during the O3 run.

In addition, commissioning time was allocated irregularly to tune or optimize some aspects of the detector, depending on the needs and opportunities. Finally, some time was spent studying and fixing problems impacting the data taking.

3.1.2. Detector steering. The Virgo data taking is largely automated and usually only requires a single operator on duty in the control room. Operators are present 24/7 during a run and take shifts every 8 hr.

The AdV detector automation, called *Metatron*, relies on the *Guardian* [37, 38] framework, developed by LIGO and based on hierarchical finite state machines. The Virgo implementation links this framework to the DAQ: automation nodes become DAQ nodes that get data directly from shared memories and are synchronized with a 1 s data availability period. A generic mechanism to read and write DAQ channels has been introduced and can be used within user codes via dedicated functions.

The full Virgo control acquisition procedure has been implemented in *Metatron*, initially prior to the O2 run and then updated for the O3 configuration, the main difference being the addition of the frequency-independent squeezing [28]. The scheme adopted, depicted in figure 4, strictly follows a top-down approach, with the lower-level nodes being automatically managed by higher-level ones.

The suspension nodes (yellow background in the graph) are tasked to align/misalign the Virgo optics, each of them is managed by the most appropriate control node (purple background), divided on the basis of the degrees of freedom to be controlled. The main node—`Interferometer Control`—is usually the only one operated manually to steer the detector. It defines the control paths, such as for instance the main global control procedure that allows reaching the Science mode (the nominal data taking state), plus other procedures to control various configurations of the optics, or to perform automated calibrations. It relies on the underlying managed nodes to perform these actions on the instrument. During the final steps of the control procedure, each single part of the interferometer is ultimately entangled with the others, and the interferometer is naturally treated as a single system. For these reasons, the last part of the procedure is directly managed by the upper level node, which sets the control parameters to the whole system, while the lower level nodes are only used as watchdogs for the correct functioning of their own sub-systems.

Additionally, the Metatron main node manages:

- The injection system, from the laser source to the IMC (Metatron node `Injection System`, orange background);
- The two OMCs, which are controlled in sequence in the final steps of the nominal control acquisition procedure (Metatron node `Output Mode Cleaner`, red background);
- The detection system at the interferometer output port (Metatron node `Detection Safety`, red background);
- The frequency-independent squeezing system (Metatron node `Squeezing System`, blue background), whose control proceeds in parallel to the one followed for the main detector. As Virgo can take valid Science data with or without this system being in its nominal state, the corresponding Metatron node is a bit apart from the others logic-wise.

Only during the calibration measurements, the `Interferometer Control` node is automatically managed by the `Calibration` node (pink background).

The Metatron framework also takes care of generating high-level flags that provide the overall status of the interferometer: this is done within the `Interferometer Status` node (green background). Finally, the `Interferometer Events` node (green background) records all state transitions of the detector. Information from these last two nodes is passed onto the Virgo live monitoring system, documented in [16].

3.1.3. DetChar organization and tools. Figure 5 shows the flow of data, from the interferometers (IFOs, on the left), to the physics analyses (on the right). While focusing on the GW candidates, this schematic highlights the three main pillars of DetChar activities during a run:

- The first timescale on which DetChar activities take place is online (latency: $\mathcal{O}(s)$). Quick automated checks are run on live data to mark out (quality: good or bad) the data stream used as input by the ‘pipelines’—that is the algorithms that scan the network data in real time, as soon as they become available. Initial data quality information is indeed shipped alongside the reconstructed GW stream, as explained in section 4.
- The second timescale is near real-time (latency: $\mathcal{O}(\text{min})$), crucial to assess the quality of the GW candidate public alerts. Thanks to a dedicated framework that is described in section 5, the data around a significant candidate are vet for each detector and a global decision is then




Figure 5. Dataflow from the interferometers (labeled ‘IFOs’ on the left) to the offline validation of GW candidates and the completion of the final dataset (right). It focuses on the generation and the vetting of the public alerts that are a key product of the LIGO-Virgo observing runs. It shows the three main timescales at which the Virgo DetChar group operates: online, near real-time and offline (see text for details).

taken: either to confirm the public alert sent to the telescopes or to retract it (see section 3.1.4 below for a description of the procedure).

- Finally, the last timescale is offline (much higher latency: up to months after the data taking). The goals of these studies are twofold: first, to finalize the dataset that all offline analyses will use, regardless of whether they look for transient or continuous signals; then, to validate the events that will be included in the final publications and whose parameters will be used to extract astrophysical information.

To ensure a continuous monitoring of the data quality, DetChar shifts were organized during the entire O3 run on a weekly basis, with two people (working onsite or remotely) on duty. The shifter crew changed every Tuesday morning, during the weekly maintenance of the Virgo detector. In addition to attending all relevant meetings, DetChar shifters usually reported their findings at the weekly DetChar meeting on Fridays and at the weekly detector meeting on Tuesdays (thus at the end of their weekly shift).

3.1.4. On call duty service and RRT meetings. An on-call service was organized during the O3 run to ensure a 24/7 expert coverage for all the Virgo detector components, from hardware systems to online computing and DetChar. In case of a problem, the operator on duty would contact the relevant experts from the control room, as well as the data taking coordinators if needed.

In addition, a joint LIGO-Virgo low-latency automated alert system was setup to contact the RRT experts—specialists of data taking, data quality or GW transient searches—who would meet remotely on short notice each time a public alert candidate was identified in real time. They would vet that candidate, using all raw information available, plus the output of several data quality checks, triggered automatically by the generation of the signal candidate: the DQR (see section 5.1 for details). The outcome of an RRT meeting could be twofold: either to confirm the public alert, or to retract it when the astrophysical origin of the candidate was questionable.

3.2. Performance

3.2.1. Noise budget. The noise budget compares the measured detector sensitivity with the incoherent sum of all known noise contributions. Each noise projection depends on the noise




Figure 6. Snapshot of the AdV O3 noise budget generated at a time of near best sensitivity of the detector (8 February 2020). The different noise sources shown are described in the text. The green line (BNS range: 66 Mpc) represents the sum of these noises: it can be compared to the measured total noise shown in black (BNS range: 59 Mpc).

level, as measured by external probes, and of its coupling to the strain channel $h(t)$, that is estimated by dedicated measurements called noise injections [32].

The AdV noise budget is based on the `SimulinkNb` [39] software package. It includes a complete model of the four main longitudinal DOFs of the interferometer (DARM, CARM, MICH, PRCL), with the interferometer optical response simulated using `Optickle` [40]. The mirror suspensions are approximated by a double pendulum state space model of the mirror and marionette (the steel body to which the mirror is suspended, a component of the Virgo suspension's last stage, called payload [41]). It also includes the feedback response measured from the transfer function between the photodiode signal and the mirror and marionette corrections. This approach allows to simply add different noise sources at their physical entry into the interferometer control loop, and also includes the expected cross couplings between the longitudinal DOFs.

This model has been verified to match the measured open loop transfer functions of the four modeled DOFs, and to reproduce the interferometer strain data calibration with errors smaller than 10%. In total, more than 100 noise sources are taken into account, and the total of those noises is summarized in figure 6.

The noises are summed in log-spaced frequency bins, which allows resolving narrow lines at low frequencies and a low statistical error on the broadband noise estimation at high frequencies. The noises taken into account are the following:

ASC – Angular Sensing and Control. This represents the control noise of 12 angular DOFs of the interferometer (two per mirror) and four DOFs of the beam injected into the interferometer. The coupling of these noises has been measured by injecting broadband noise into each DOF [42].

DAC – Digital Analog Converter. This is the electronic noise of the digital to analog converters used to drive the six main mirrors and marionettes of the interferometer. This electronic

noise has been measured in the laboratory before installation, and the noise coupling is modeled using `SimulinkNb`.

Dark. This is the electronic and dark noise of the photodiodes used in the four longitudinal DOFs control. The dark noise of each photodiode is measured by closing the mechanical shutter in front of it. The noise coupling is part of the model of the longitudinal control loops in `SimulinkNb`.

Demodulation. This is the phase noise of the demodulation of radio-frequency signals from photodiodes to control CARM, MICH and PRCL. That phase noise mixes the two demodulation quadratures. This bi-linear noise source is measured, and the noise coupling is modeled using `SimulinkNb`.

ENV – Environment. This is the sum of three contributions: acoustic, magnetic and scattered light. The acoustic and magnetic noises are measured with four microphones and three-axis magnetometers, located in the experimental buildings near the interferometer components (see [32, 33] for details). Their couplings are measured by broadband and sweeping sine noise injections. Scattered light noise is projected in two ways: (i) using the measured relative intensity noise on auxiliary photodiodes and measuring a linear coupling by shaking the bench hosting that photodiode to elevate the noise in the detector sensitivity; (ii) using position sensors of suspended benches that couple in a non-linear way, with a modeled coupling that is scaled based on measurement obtained by displacing intentionally the bench by tens of microns per second to elevate the noise in the detector sensitivity [43].

LSC – Length Sensing and Control. This represents the control noise of four DOFs: MICH, PRCL, OMC length, and residual intensity noise. The noise is measured in all cases, the coupling is measured for all except for the OMC length where it is modeled. Note that this results in double counting the dark and quantum noise of the sensors used for MICH and PRCL control, however these double counted contributions are negligible.

Quantum. Quantum noise of the detector and shot noise of the sensors used for MICH, PRCL and CARM control. The noise and the coupling are modeled using `SimulinkNb`.

SSFS. This represents the control noise of the relative error between CARM and the laser wavelength. The noise is measured, the frequency dependent coupling is modeled using `SimulinkNb` and a time dependent scaling factor is measured.

Seismic-Thermal. This is the sum of the negligible seismic noise and three thermal noise contributions: suspension, mirror coatings and residual gas pressure in the arm vacuum tubes. The noise sources and the couplings are modeled using analytical functions in separate dedicated codes.

‘flat noise’. It is a noise source of not yet understood physical origin. Its level has been measured proportional to the square root of the DARM offset used to obtain the interferometer DC readout [44, 45].

The sum of the noises described above correspond to a BNS range of 66 Mpc, while the actual BNS range in the corresponding data was measured at 59 Mpc. Hence, about 10% of the noise limiting BNS detections is unaccounted for, not understood and not described in this section.

More in detail, at frequencies above 1 kHz the sensitivity is mostly limited by quantum shot noise. The measured level is about 5% higher than expected. This is due to a slow degradation of the frequency-independent light squeezing during O3, from 3 dB at the beginning of the run to about 2.5 dB at the end of it.

In the most sensitive frequency range, between 80 Hz and 200 Hz, there are significant contributions from three sources: quantum shot noise, mirror coating thermal noise and the “flat noise” of unknown physical origin. Assuming that the ‘flat noise’ estimate is correct,

Table 1. Summary of the Virgo global control acquisition performance during O3: the control is acquired after a successful control acquisition *sequence* that counts one or more control acquisition *attempts*.

Global control acquisition attempt	
Median duration	18 min
Distribution of this time	
Reaching the detector working point	~30%
Controlling the two OMCs	~50%
Acquiring the lowest noise configuration	~20%
Global control acquisition sequence	
Median number of attempts	2
Median duration	25 min

removing completely this unknown noise source would have resulted in 10 Mpc improvement in the BNS range.

At low frequencies between 20 Hz and 50 Hz, the dominant noise sources are quantum radiation pressure noise that is increased by the frequency independent light squeezing and the laser intensity noise. However, 30% of the noise remains not understood in that frequency range, so other significant noise sources are yet to be identified.

3.2.2. Virgo O3 duty cycle. Table 1 summarizes the performance of the global control acquisition procedure for the Virgo detector during O3. This performance has been stable over the whole run, showing the robustness of that procedure. As not all control acquisition *attempts* are successful, a global control acquisition *sequence* is defined as a set of successive control attempts that leads to the global control of the instrument.

The dataset analyzed here spans the whole O3 run and includes more than 700 successful global control acquisition sequences. Only the periods during which detector activities incompatible with data taking (maintenance, commissioning, calibration and known hardware problems) were ongoing have been excluded. In order to be less sensitive to the tails of the statistical distributions—which do impact the duty cycle (see the ‘Locking’ contributions to the pie charts below) but can have multiple origins (human errors, hardware or software failures possibly hard to diagnose quickly, or external conditions like bad weather) which are not directly related to the global control acquisition procedure—we have decided to report median durations in the following.

The median duration of a successful global control acquisition attempt is 18 min: 30% of this time is spent reaching the detector working point (Michelson interferometer at the dark fringe, power recycling cavity and arm cavities resonant, SSFS enabled); 50% is spent to control the two OMCs at the Virgo output port; the final 20% are used to reach the lowest noise configuration at the level of the suspension actuation. The median number of attempts needed to complete a global control sequence is 2 and the median duration of a successful global control acquisition sequence is 25 min. During O3 the quickest sequence took however 13 min.

Table 2 details the control stability of the Virgo detector, separately for the sub-runs O3a and O3b, and averaged over the whole O3 run. The ‘global control segments’ are stretches of data during which Virgo is controlled in its nominal low-noise configuration, while, as already defined, the ‘Science segments’ are the subset of the global control segments during which Virgo is taking data of good quality, to be used by analyses. The difference of duration between

Table 2. Summary of the O3 data taking performance of the Virgo detector. The last three rows of the table provide duty cycles for different configurations of the 3-detector LIGO-Virgo global network: the fraction of the time during which at least one the three instruments is taking data, at least two are and finally all three are.

		O3a	O3b	O3
Virgo global control segments	Mean [hr]	6.1	6.4	6.3
	Median [hr]	2.7	1.8	2.2
Virgo Science segments	Mean [hr]	5.0	4.0	4.5
	Median [hr]	2.6	1.4	1.9
Duty cycles	Virgo [%]	76.3	75.6	76.0
	Network—at least 1/3 [%]	96.8	96.6	96.7
	Network—at least 2/3 [%]	81.9	85.4	83.4
	Network — 3/3 [%]	44.5	51.0	47.4

the global control and Science segments is dominated by limited disruptions of the data taking, that usually stop the Science mode for a short time. The dominant source of these breaks is the frequency-independent squeezer that lost its nominal configuration about 240 times during the O3 run; the median time to restore it and switch back to Science data taking was about 140 s.

We note that the Virgo segment duration summary numbers listed here are lower than those reported by LIGO [46, 47]. Yet, this difference has no significant impact on the duty cycle that is very similar for the three detectors of the global LIGO-Virgo network. The comparison between the O3a and O3b sub-runs shows that the impact of the winter season (larger sea seismic activity, wind, and more generally bad weather), although real, has been limited. Overall, the global network duty cycle has improved during O3, mainly due to the increase of the LIGO detectors duty cycle, while the Virgo one has been very stable. With an average of 76%, the Virgo O3 duty cycle is lower than that measured during August 2017, the final weeks of the O2 run Virgo took part of: $\sim 85\%$. Yet, the O3 performance has been achieved over 11 months spanning a whole calendar year and cannot be directly compared to the duty cycle of a short (only 25 days) run in Summer time, the most favorable period to operate an instrument like Virgo. Running one full year instead of one month is also more complex person-power wise, and the Virgo organization implemented during O3, although perfectible, held on during the whole run. This experience represents a good base on which to build upon in order to improve the Virgo performance for the O4 run and beyond.

Figure 7 shows the breakdown of the time spent in different modes by Virgo during O3. Overall, the O3a and O3b distributions are quite consistent. Breaking these 11 month-averaged duty cycle figures down to a 24 hr period, Virgo took data during 18 hr, with the remaining 6 hr roughly divided into three blocks of the same duration: ~ 2 hr for controlling the detector (Locking), ~ 2 hr for recurring activities (Calibration, Commissioning and Maintenance) and ~ 2 hr for solving issues (Any other state).

The analysis of these pie charts shows that increasing the duty cycle during future runs will not be straightforward. The room for improvement is limited in each area and so any significant duty cycle gain will likely stem from a combination of various small progresses, each made possible by the redesign or the optimization of a particular process.

To conclude this overview, figure 8 summarizes the improvement of the sensitivity of the AdV detector. The BNS range associated to each curve is given in the legend. From O2 to O3b, the record BNS range has more than doubled from 28 Mpc to 60 Mpc, with a continuous




Figure 7. Breakdown of the time spent in different modes by the Virgo detector during the O3 run (larger, top-middle pie chart) and separately during the O3a and O3b sub-runs (pie charts at the bottom). The ‘*Locking*’ mode corresponds to periods when the control of the detector is being acquired. The regular *maintenance* and *calibration* periods have been described in section 3.1.1. Finally, the *Any other state* category includes all the other situations encountered during the whole run: troubleshooting periods, various kinds of tuning, etc. These results exclude the 1 month-long commissioning break that took place in October 2019, in between the O3a and O3b sub-runs. In each pie chart, the modes are sorted by decreasing percentage.




Figure 8. Comparison between four sensitivity curves of the AdV detector: during O2 (green trace), at the beginning of O3a (purple), at the beginning of O3b (red) and at the end of O3b (blue). The caption provides the corresponding estimated BNS ranges.

improvement of the sensitivity in the whole bandwidth of the detector. Many spectral features of the residual noise structures have either been removed or significantly reduced over time.

3.2.3. The Virgo O3 dataset. The final Virgo O3 dataset consists of more than 250 days of data recorded during the O3a and O3b sub-runs and whose quality has been checked and validated (described in section 6.4.1). It is built upon and supersedes the online good-quality Science dataset that was used as input by the analysis pipelines that looked for GWs in real time (see section 4). Dedicated studies have been performed offline to refine the quality assessment of the data. In addition to running more in-depth analyses, new checks have been added during the run, as potential flaws got discovered in the existing analyses, or new problems identified at the detector level. Moreover, small sets of good data that had not been automatically included in the dataset (either because they were incorrectly labeled or because part of their data quality information was missing) were added by hand.

The main categories of checks applied to assess the quality of the Virgo data are the following:

- **Are key components of the Virgo hardware (suspensions and photodiodes) having transient problems?**

These checks, described in section 4.1.2, were fast enough to be performed online on live data.

- **Is the reconstruction of the GW strain time series $h(t)$ nominal?**

This is a prerequisite for any further use of the Virgo data. The online reconstruction of the Virgo data was satisfactory during O3: only data from the very end of O3a (16–30 September 2019) were reprocessed offline to increase the sensitivity by a few percents [34]. Yet, during periods of high seismic activities (bad weather, high wind or the passing of seismic waves from strong and distant earthquakes), the nominal global control configuration could be replaced [33] by a more robust one, the so-called ‘earthquake (EQ)-mode’ [34]. Although that procedure saved some losses of the working point (whose recovery would have costed time), it could not be validated against the nominal reconstruction of the $h(t)$ strain stream until the final two months of O3b. Therefore, during most of the O3 run, data taken in these peculiar conditions had to be excluded from the final dataset.

- **Do the data suffer from known problems?**

Tailored checks were run offline to identify and isolate periods during which the detector was not behaving nominally, although it was still controlled. One example of such studies is the fact that the NI mirror suspension was randomly suffering from transient (a few second-long) losses of data. This was usually enough to lose the control of the entire detector, and hence to lose at least 20 min to 30 min of data: the time to reacquire the global working point and to restore Science data taking. Therefore, a patch was developed by experts to detect the data loss and switch in real time to a less robust—but still available—control until the missing data were back. This saved hours of running time for Virgo overall, but a dedicated scan of the data had to be performed offline to identify the occurrences of these control switches (potentially inducing transients and artifacts of instrumental origin in the data) and to remove them from the final dataset.

- **Are the data consistent?**

The last few seconds of a segment preceding a control loss of the detector have been removed offline, as the $h(t)$ data could be corrupted (see section 6.4.1 for details). In addition, we have verified that the detector was nominally controlled during all segments flagged online




Figure 9. Distributions of the Virgo BNS range during O3 for the online (red histogram and trace) and offline (blue) dataset. The top (bottom) plot compares the histograms (cumulative distributions).

as Science, and we also have looked for segments which could be included in the final offline dataset although they had not been categorized as Science online.

- **Is the dataset complete?**

There could be data segments with missing or corrupted $h(t)$ channel that would require a limited reprocessing. Or there could be segments with missing data segments due to problems in the DAQ, etc. Dedicated checks were setup to target these problems specifically, before analysts would run into them when processing the data.

Data segments that fail one of the checks defined above are classified as ‘Category 1’ (CAT1) vetoes and are excluded from all analyses. Overall, only 0.18% of the Virgo O3 Science dataset have been CAT1-vetoed.

To conclude this overview of the Virgo performance during the O3 run, figure 9 compares the Virgo BNS range distributions before (red) and after (blue) applying data quality cuts to determine the final O3 dataset. As expected, data quality requirements remove periods of low BNS range, i.e. when the sensitivity was poor. Yet, about 1% of the data have a BNS range lower than 35 Mpc, that is significantly below the typical values achieved during O3 for that sensitivity estimator. While these data have not been flagged as bad by the various checks run on the dataset, they correspond to periods during which the detector was less accurately controlled, in particular due to bad weather.

4. Real-time data quality

Online data quality was a key challenge to tackle for DetChar during the O3 run. The availability and the reliability of that information, supporting the data taking, had to be high in order to allow the real-time transient GW searches to make the best use of the Virgo data. Significant candidates identified by those analyses—usually found in data from at least two of the three detectors of the global network, but sometimes identified in a single instrument—would then




Figure 10. Online architecture to produce data quality products during the O3 run. The status of the interferometer is monitored by a dedicated Metatron server (see section 3.1.2). Data quality flags are generated by dedicated servers described in [16]: the DMS, the BRMSMon process (environment), the VetoMerger process (large deviations in auxiliary signals), the Hrec process ($h(t)$ reconstruction), and the Omicron algorithm (glitches in $h(t)$ and auxiliary signals). Data quality segments are then generated by the SegOnline process and saved in the LIGO-Virgo segment database, while the online $h(t)$ stream, the state vector and the veto channels are sent to online data analysis pipelines through the V1FromOnline server.

lead to public alerts, used by telescopes worldwide to search for counterparts of potential GW signals.

In this section, we first describe the different blocks of the Virgo online data quality architecture, in use at EGO during the O3 run. This framework matches the dataflow shown in figure 5 and is complemented by the vetting of the most significant triggers identified in low latency, described in the following section 5.

Real-time information about the detector status was combined with fast data quality estimators to produce a single integer channel sampled at 1 Hz, the *Virgo state vector*. That state vector was shipped alongside the GW strain channel $h(t)$ to computing centers where data were analyzed in real time. Its integer value was constructed by gathering several binary information (schematically: good vs. bad) encoded as bits; that bit pattern would later be decoded by the analysis frameworks to discard any bad data. Parallel to this data analysis stream, this information—the detector status plus the real-time assessment of the data quality—was automatically uploaded by a dedicated online process (called SegOnline) to the Data Quality Segment Database (DQSEGDB) [16].

Finally, we present the experience gained during O3 with additional data-quality inputs, called *veto streams* whose aim is to help searches to reduce their false alarm rate by identifying triggers that are very unlikely to be of astrophysical origin.

4.1. The Virgo O3 online data quality framework

The online data quality architecture is designed to deliver data quality products to online transient searches. It is based on a set of servers connected to the DAQ and providing relevant information about the quality of the data (the whole raw data, plus the reconstructed $h(t)$ stream). In the following, the main elements of this architecture, summarized in figure 10, are presented.

Table 3. Definition of the bits of the Virgo state vector during the O3 run (see text for details).

Bit number	Active when
0	$h(t)$ successfully computed.
1–2	Science mode enabled.
3	$h(t)$ successfully produced by the calibration pipeline.
4–7	Bits irrelevant for the present discussion: either redundant with other bits or unused during O3.
8	No DetChar-related hardware injection (see section 6.2 for more details).
9	No continuous wave hardware injection (the only type of non calibration-related injections performed for a short period during O3, while taking nominal data).
10	Online data quality is good (no CAT1-type veto).
11	Virgo interferometer fully controlled, with a nominal working point or close to it.
12–15	Not used.

4.1.1. State vector. Table 3 defines the 16 bits of the Virgo *state vector* integer channel in use during the O3 run. A bit is said to be *active* when its value is 1, meaning that the corresponding check is passed. A value at 0 means instead that a problem, or a non-nominal state, has been detected. The information provided by these bits is on purpose partially redundant, in the sense that several bits can be at 0 when proper data taking conditions are not met. During O3, the bits 0, 1 and 10 were required to be active to have the corresponding 1-second data frame processed by real-time analyses.

4.1.2. Online CAT1 vetoes. During the O3 run, the problems detected online and leading to CAT1 vetoes are listed below. These are:

- No saturation of any of the 4 dark fringe photodiodes, using the ‘DC’ (from 0 to a few Hz) and ‘Audio’ (from a few Hz to 10–50 kHz) demodulated signals.
- No saturation of the correction signal of any of the 16 suspension stages monitored.
- No saturation of the rate of glitches reported by the online `Omicron` framework for the DARM correction channel¹¹⁶.

These saturation checks were combined using a logical OR to produce CAT1 vetoes with a 1 s granularity. Section 6.4.1 describes the corresponding set of *offline* CAT1 vetoes, used by all analyses processing the final O3 Virgo dataset, including these *online* CAT1 vetoes.

4.1.3. SegOnline. Any channel provided by the DAQ or by the online processing (for instance DMS monitors or the BRMSMon process) can be used by the `SegOnline` process to build segments of data quality flags which are sent online to DQSEGDB [16].

¹¹⁶ A more correct way to monitor the glitch rate would have been to scan $h(t)$, but the latency added by that `Omicron` check would have delayed too much the GW strain channel for online processing. The offline equivalent version of that check did use $h(t)$, as latency was not an issue anymore in that case.

SegOnline writes down segments into XML files with a latency of about 10 s and those XML files are then read by a *rsync* process to upload the segments into DQSEGDB every 5 min. Such data quality segments can then be used by any analysis, or can be viewed and downloaded through a dedicated web interface [48].

4.2. Veto streams

Low-latency transient searches are limited by glitches in the $h(t)$ data. Each search pipeline is sensitive to specific families of glitches. The online data quality architecture is designed to deliver a channel to flag glitches relevant to a given low-latency pipeline. These channels are called veto streams. A veto stream is a time series which can only take two values: 0 means good quality and 1 bad quality. A veto stream is generated by the VetoMerger process which combines information from many online data quality processes, carefully selected to target the glitches limiting the search of interest [16].

Some Omicron processes [16] are configured to select triggers detected in auxiliary channels with a SNR above a threshold tuned with the UPV algorithm [16]. These triggers are known to witness glitches in the $h(t)$ channel. When this is the case, the veto channel is set to 1. VetoMerger also ingests the data quality flags generated by BRMSMon [16] to veto environmental disturbances.

In O3, the veto stream system was experimented as an input to one of the low-latency searches for compact binary mergers, PyCBC Live [49, 50]. The veto stream, named DQ_VETO_PYCBC, combined two elements: a veto channel delivered by Omicron to target scattered-light glitches, and a data quality flag produced by BRMSMon to tag occasional glitches associated to lightning strikes. As explained in section 6.2, it is critical that a veto is constructed from channels which are insensitive to GWs: a channel (or veto) is then said to be *safe*. The channel safety is tested with hardware injections that mimic the effects of GWs on the detector. The DQ_VETO_PYCBC veto stream is derived from safe channels: magnetometer signals are used to veto lightning strikes and DAQ frequency modulation channels were found to witness scattered-light glitches. A conservative approach was adopted to tune the vetoes: their thresholds were set at high values to reliably flag really limiting glitches, while keeping the rejected time low. As a result, only 0.05% of the O3 Science time was flagged by the DQ_VETO_PYCBC veto stream. This is a conservative tuning as it means that there is only a 0.05% probability that DQ_VETO_PYCBC would discard a true GW signal, under the reasonable assumption that GWs are uncorrelated with scattered-light and lightning glitches. PyCBC Live used the veto stream to simply prevent the generation of a candidate event from Virgo data, or remove Virgo's contribution from a LIGO-Virgo candidate, during periods of active veto. In future runs, the veto streams may be integrated in a more general framework based on auxiliary channels to discard or down-weight transient noise events.

Although the veto stream was only used for the online PyCBC analysis in O3, we can now evaluate its effect on the PyCBC offline analysis as well. In particular, we consider here Virgo single-detector triggers generated by the broad-space PyCBC search [51] during the period from 1 April to 11 May 2019. The study is performed under the assumption that such triggers are dominated by noise, and that scattered-light and lightning glitches are not correlated with GW signals. The search ranks the triggers by a quantity known as *reweighted SNR* [51], i.e. the SNR returned by the matched filtering technique [52, 53], weighted by the result of χ^2 tests that quantify how well the time-frequency distribution of power observed in the data is consistent with the one expected from the matching template [54, 55]. For practical reasons, only triggers with reweighted SNR higher than 6 are considered here. After this selection, the sample is composed of roughly 2.5×10^5 triggers. To evaluate the impact of the vetoes on the offline




Figure 11. Effect of veto streams on triggers from the O3 archival search for compact binary mergers in Virgo data based on PyCBC. Each staircase shows the cumulative fraction of vetoed triggers with reweighted SNR higher than a given threshold, as a function of the threshold. The red curve shows what happens when removing triggers within segments flagged by the veto streams. The gray curves, instead, are samples from the null distribution, constructed by applying a time shift to the segments before vetoing the triggers. The fractions are relative to the overall number of triggers generated by the search. The left plot considers vetoes targeting scattered-light glitches, and the right plot considers vetoes associated with glitches from lightning. In both cases we can see that the red curve is significantly above the gray ones, at least for some reweighted SNR thresholds, indicating a significant correlation between the triggers and the vetoes.

search, we remove triggers with a merger time belonging to a vetoed segment. We carry out the study separately for vetoes targeting scattered-light glitches and glitches from lightning. We show the fraction of vetoed triggers as red staircases in figure 11, and in both cases it is found to be of the order of 10^{-4} or less, i.e. compatible with the expectation from the amount of vetoed time. The majority of vetoed triggers have relatively small reweighted SNRs.

As a next step, we would like to assess the statistical significance of the impact of the vetoes on the offline PyCBC triggers, i.e. calculate the probability that the vetoed triggers are simply explained by chance alignment with the veto segments. To this end, we shift the veto segments rigidly by a constant time offset, so that the shifted segments maintain the general time structure of the original ones, but lose any correlation with glitches. We then recompute the fraction of triggers vetoed by the shifted segments, obtaining a null sample that we can compare to the original fraction from the unshifted segments. Note that the null fraction is rescaled to account for the overlap between the Science-mode segments and the time-shifted veto segments, which is a function of the time shift. We construct 1000 such null samples by repeating the time-shifted analysis with time offsets covering the range of $[-50\,000, +50\,000]$ s in steps of 100 s. The cumulative fraction of vetoed triggers for the null samples are shown in gray in figure 11. At a reweighted-SNR threshold of 6, the unshifted fraction is higher than any time-shifted fraction, for both scattered-light (left plot) and lightning (right plot) vetoes. We conclude that the probability for the observed effect of the vetoes on the PyCBC offline triggers to be a statistical fluctuation is less than the inverse size of our null sample, i.e. 10^{-3} . For higher reweighted-SNR thresholds of 8.5 (10.5), this probability is 2×10^{-3} (2×10^{-2}) for vetoes targeting scattered-light glitches, and less than 10^{-3} (6×10^{-3}) for vetoes associated

with lightning glitches. It does therefore appear that scattered light and lightning strikes are correlated with a small population of (relatively weak) PyCBC triggers, and that the veto streams can in principle be used to remove or down-weight these triggers. Whether this is beneficial or not should be established by carrying out a complete simulation with injected signals, in order to measure the effect of the veto streams on the sensitive time-volume of the search, as routinely done when optimizing searches for compact binary mergers. We reserve this to a more detailed future investigation, however our results indicate that this effect would not be very large (assuming O3-like data) as the fraction of vetoed signals is of order 10^{-4} and the reduction of the background would only impact relatively quiet Virgo triggers anyway.

For the O3 run we demonstrated the possibility of delivering search-specific veto streams to online pipelines to reject transient noise events. This first experiment with PyCBC, although with limited performance, has validated the online architecture. For the next run, O4, we plan to generalize this framework to other pipelines and plug in more veto streams to target other families of glitches.

5. Public alerts

As demonstrated with the extraordinary GW170817 [5] event from the O2 run, public alerts sent by the LIGO-Virgo network are key deliverables targeting the astronomy community. Yet, how successful these are depends on the accuracy of the information provided, and of the latency at which they are delivered. For O3, the main contribution of the DetChar group to this effort has been the design and the implementation of the DQR framework. A DQR is a set of data quality checks, automatically triggered by the finding of a new GW candidate. Its output allowed the RRT team to vet the associated data in a timely way. Moreover, its usage extended way beyond the data taking period, as it was the main tool used to assess the data quality of all GW candidates identified by analyses, in some cases with a latency longer than a year (compared to when the corresponding data were acquired).

The performance of the Virgo O3 DQR is described below, before summarizing how Virgo contributed to the LIGO-Virgo public alerts during the O3 run. The Virgo DQR implementation can be found in [16].

5.1. Performance of the Virgo O3 DQR framework

This section briefly summarizes the performance of the Virgo DQR, via statistical analyses using data from O3b that correspond to the final, most complete, version of that framework during the O3 run. Emphasis is put on latencies and running times, as these are key quantities to vet public alerts in a timely way. As those time distributions can include tails due to occasional technical problems impacting the DQR dataflow somewhere along its way (from the GraceDB) [16] to the EGO HTCondor farm and back) while being external to it, the results presented in the following two tables include the 50th and 95th percentiles in addition to the mean values.

Table 4 provides the measured latencies for the processing steps that occur upstream of the DQR with the following meaning:

- The first line is the difference between the time when the trigger is recorded in GraceDB and the time when the corresponding data were acquired.
- The second measures the time needed for GraceDB to notify the LVA1ert and to have this message trigger the Virgo DQR framework upon reception.
- The third reports the time needed to create and configure a new DQR instance, until it is ready for processing. One should note that this duration includes a 300 s wait time, imposed

Table 4. Summary of the performance of the low-latency with Virgo DQR dataflow during O3b, from the GPS time of a trigger to the start of the Virgo DQR on the EGO HTCondor farm: see text for details.

Operation	Time taken (s)		
	Median	Mean	95th percentile
Data acquired → Candidate on GraceDB	52	166	331
Candidate on GraceDB → LVA1ert trigger	4	4	11
LVA1ert trigger → Virgo DQR configured	331	339	383
Virgo DQR configured → Virgo DQR started	8	10	21

Table 5. Summary performance of the Virgo DQR processing during the last ~ 100 days of the O3b run. The quoted durations include the time to upload DQR check results back to GraceDB that usually takes from ~ 5 to ~ 20 s.

Operation	Time from DQR start (s)		
	Median	Mean	95th percentile
Quick key checks	374	383	619
Adding Omicron trigger distributions	868	816	935
Adding full Omicron scans	1740	2159	4690
End	5185	4954	6330

in order to allow GraceDB to receive, process and gather all triggers found by the different online searches that analyze strain data in parallel and independently. The assumption is that, after these 5 min, the low-latency information available in GraceDB should be optimal and stable in the vast majority of cases. Therefore, the actual DQR configuration phase only takes a few tens of seconds: the needed data are located in the low-latency streams just made available by the DAQ and more than 30 scripts (each corresponding to a data quality check) are generated one after the other.

- Finally, the last reported duration accounts for the time needed to start processing the DQR on the EGO HTCondor farm. This depends on the occupancy of the farm and of the EGO internal network performance.

We can see that the mean time elapsed between the recording of the data by the different detectors and the creation of a new record in GraceDB is under 3 min. The median time is even under 1 min while the tail of the time distribution extends beyond 5 min. This includes the reconstruction of the GW strain channels, the transfer of these data alongside the associated online data quality information to computing centers, the processing of these data by real-time GW searches, the automated analysis of the results and the final transfer of trigger information to GraceDB where it is made centrally available. Then, the new alert is received at EGO a few seconds later, triggering the creation and the configuration of a new DQR instance. Removing the compulsory wait time of 300 s, the DQR configuration takes a few tens of seconds only. Finally, 10 additional seconds are needed on average to have the first DQR jobs processed on the EGO HTCondor farm.

Table 5 summarizes the performance of the Virgo O3 DQR framework in terms of running time. Each row corresponds to a category of checks. The quoted durations increase from one row to the next, as each new set of checks includes the previous ones.

Table 6. Percentages of the O3b Virgo DQRs with 0, 1 and 2 unsuccessful checks respectively.

Number of unsuccessful checks	0	1	2
Percentage of O3b automatically processed DQRs	85%	13%	2%

The quick checks whose outputs are mandatory to vet a trigger take about 6 min to be all available, with a few minutes spread. Adding information about the *Omicron* triggers around the candidate takes about 10 more minutes. During O3, this latency was dominated by the fact that *Omicron* triggers were computed in real time and stored internally by the online server: they were only written to disk every 600 s, in order to allow the framework to cope with the incoming data flow. Work will be done prior to O4 to optimize this latency and to make the DQR aware of when the needed data have been written to disk, so that their processing can start immediately after. *Omicron*-scanning all the available channels (more than 2000 in total, with the vast majority of them sampled at 10 kHz) around the trigger time requires 15–20 additional minutes. Finally, the full DQR took from 1.5 to 2 h to complete. The longest checks were BruCo [16] and UPV, plus the scan of all online logfiles described above.

The Virgo DQR reliability—that is how efficient a DQR is in providing accurate information on a GW trigger—is another key figure of merit of that framework. A good indicator of this is the number of (software) check failures per DQR instance, as any check not properly completing would prevent analysts from accessing part of the available data quality information. During the whole O3 run, there was no case of a public alert for which the rapid vetting of the Virgo data was delayed due to Virgo DQR issues. In addition, we used the large sample of DQRs automatically processed during O3b—by design, all GW candidates with a low-latency false-alarm rate below 1/day triggered a DQR. The results are in table 6: only 13% (2%) of the DQR had 1 (2) failed checks—and none had more than 2 failures, while an O3 DQR included more than 30 checks in total. No exhaustive analysis of these failures has been performed, as most of these DQRs were never checked by hand because the associated trigger was not significant enough. The two main causes of problems were, however, incomplete handling of edge-cases with the input data and actual bugs in DQR check algorithms. These issues are being addressed as part of the upgrade of the Virgo DQR framework for the O4 run.

5.2. O3 public alerts

5.2.1. Public alerts retracted because of an issue with Virgo data. During O3, 24 public alerts out of 80 have been retracted: 8 during O3a and 16 during O3b. Out of these retractions, only two were due to Virgo data:

- S191124be [56] was due to a problem in the noise removal procedure included in the online reconstruction of the $h(t)$ GW stream [34]. Two such cleaning algorithms running in sequence started interfering, leading to a noise increase over time. An online pipeline started triggering on that excess noise, creating several non-astrophysical GW candidates in rapid succession (figure 12), until one of them had a false alarm rate lower than the public alert threshold. That led to the generation of an automated alert that was then quickly retracted, the problem of the noise cleaning procedure was fixed as well.

A similar problem should not happen again in future runs for three reasons: (i) improved noise cleaning procedures are being developed within the Virgo $h(t)$ reconstruction; (ii) an online monitoring dedicated to such noise removal interferences will be in place during O4; (iii) a monitoring of the pipeline trigger rates in GraceDB will be running as well during




Figure 12. Number of Virgo DQRs automatically processed per day during the O3b run. The peak of 31 entries corresponds to 24 November 2019 when there was a transient problem with the Virgo $h(t)$ reconstruction: that generated several online triggers, finally including S191124be that passed the public alert threshold and was promptly retracted.

future data taking periods, in order to spot quickly any misbehavior, like an excess trigger rate (in the case of S191124be) or the opposite: a too long data-taking time period without any trigger, even of low significance.

- S200303ba [57] was a single-pipeline trigger with most of its SNR concentrated in Virgo. At that time, Virgo data were quite noisy due to bad weather. An Omicron-scan around the trigger time (figure 13) showed evidence of scattering light noise at low frequency. The unusually-long delay to send out the retraction circular (80 min) was partly due to an issue with the Gamma-ray Coordinates Network broker connection.

5.2.2. Virgo contribution to O3 public alerts. Out of the 56 non-retracted O3 public alerts, 42 involved the Virgo detector. For 10 out of the 14 LIGO-only alerts, Virgo was not controlled in its nominal configuration at the GPS time of the trigger. This fraction is consistent with the average duty cycle of Virgo during O3 (see section 3.2.2). For the four remaining alerts, described in detail next, Virgo was fully controlled at the time of the trigger and had a BNS range consistent with its typical performance at that moment.

S190720a occurred during a ~ 1 min segment in between the end of the acquisition of the detector global working point and the beginning of the nominal Science mode. Therefore, Virgo data were not used for low-latency analyses. Offline analyses later confirmed S190720a as a significant detection and were able to use the low-noise Virgo data, finding a non-negligible amount of signal power in them. S190720a was published as GW190720_00 0836 in GWTC-2 [8].

S190910d occurred during nominal Science mode data-taking in Virgo. It was a marginal candidate, only reported by a subset of the low-latency searches. These searches did not find a significant amount of signal power in Virgo data, and did not report Virgo as being used for the candidate. S190910d was not confirmed by offline analyses.




Figure 13. Omicron spectrogram around the time of the S200303ba trigger. The search template time-frequency track (solid line) overlaps with low-frequency scattering-light glitches (yellow color) caused by bad weather.

S190923y occurred while Virgo was undergoing commissioning activity. It was not confirmed by offline analyses.

S200225q occurred while Virgo was undergoing a calibration run. Offline analyses confirmed S200225q as a significant detection and were able to include the low-noise Virgo data, although no significant signal power was found there. S200225q was published as GW200225_06 0421 in GWTC-3 [10].

6. Global data quality studies

This final section presents examples of global data quality studies made during or after the O3 run: noise transients, spectral analyses, classification of auxiliary channels based on their potential sensitivity to GW signals and offline data quality studies leading to the final Virgo O3 dataset.

6.1. Glitches and pipeline triggers

6.1.1. Glitch rates during the O3 run. During data taking, Omicron runs online on a few hundred channels, including the GW strain $h(t)$, and monitors glitches in real time: these triggers are stored on disk with a few minutes latency. Figure 14 displays the evolution of the glitch rate during the O3 run, and constitutes a reference for the detailed view of figure 15, where the global Omicron glitch rate has been broken down into SNR (top plot) and peak frequency (bottom plot) bands. In these plots, the rates have been averaged by computing their daily moving median to ease the reading. From the top plot in figure 15, we can notice that the ratio




Figure 14. Virgo glitch rate, using Omicron triggers, for the final O3 dataset (Science segments that have not been CAT1-vetoed). The blue dots are averages over one hour while the red curve shows the corresponding daily moving average. The gap in between O3a and O3b corresponds to the 1 month commissioning break.

of glitches in the various SNR bands is approximately constant. On the contrary, the bottom plots highlight several temporary increases in glitch rates, different for the various frequency bands. The choice of these frequency bands has been made to try to isolate some possible classes of glitches, and consequently of their sources: the region below 45 Hz is characteristic of scattered-light glitches, enhanced during bad weather conditions; the regions around 50 and 150 Hz contain the mains fundamental frequency in Europe and its second harmonics, hence is likely related to glitches of electrical origin; the region around 450 Hz contains the frequencies of the suspension wire violin modes for the test masses, and of another harmonic of the mains frequency.

The large majority of glitches identified by Omicron have a moderate SNR: between 5 (the minimum value from which the Omicron trigger is kept) and 8. The highest trigger rate at the very beginning of O3a corresponding to glitches with a peak frequency between 440 Hz and 460 Hz is an artefact due to a mis-configuration of the Omicron online server that was quickly fixed. The significant increase of the trigger rate in O3b with respect to O3a is mainly due to the bad weather conditions during the fall and winter seasons (see [33] for more details). The weather was actually very quiet in January 2020 and the associated drop in glitch rate is quite strong.

6.1.2. Offline searches. Non-stationary instrumental noise can potentially impact searches for transient GWs, which must include methods to robustly separate astrophysical candidates from noise fluctuations. Despite the power of such methods, inspecting the candidates produced by a search remains an important way to identify problematic operating conditions of GW detectors, and to understand if the search needs to be tuned in different ways for different detectors.




Figure 15. Glitch rates (daily moving average) using Omicron triggers during the O3 run for Virgo. The gap in between O3a and O3b corresponds to the 1 month commissioning break. The top plot breaks the glitch rate into SNR ranges, while the bottom one categorizes it in terms of frequency ranges for the glitch peak frequency.

In this section, we perform a detailed inspection of the candidates produced from Virgo O3 final data by one of the analyses used for compiling the GWTC catalog, namely the PyCBC offline analysis [51] which we already considered in section 4.2. As a reminder, this analysis performs a broad-space search for compact binary mergers involving neutron stars, black holes, or both. It uses a bank of model waveforms and matched filtering to generate candidates from LIGO and Virgo data. Each single-detector candidate is ranked by a reweighted SNR, i.e. a combination of its matched-filter SNR and two χ^2 signal-based discriminators [54, 55] designed to reject candidates produced by non-stationary noise. Such discriminators have been




Figure 16. Rate of compact binary merger candidates produced from Virgo O3 data using a broad-space search based on PyCBC. Left: rate as a function of the matched-filter SNR. Right: rate as a function of the reweighted SNR, which combines the SNR and two χ^2 discriminators. Red, blue and green curves correspond respectively to templates lasting less than 0.6 s, 0.6–4 s, and more than 4 s. Curves of the same color represent different chunks of Virgo data, each lasting ~ 5 days. The black curves show the rate over the entire run and search space. They extend to lower rates than the individual chunks due to the much longer duration of the entire run.

mainly tuned on noise from Advanced LIGO so far, and to the best of our knowledge, their behavior on Virgo data has not been published before.

Figure 16 shows the rate of candidate events recorded by PyCBC from Virgo data. The horizontal axis shows either the matched-filter SNR of the candidate (left plot) or its reweighted SNR (right plot). The vertical axis shows the rate of candidates that are ranked higher than the value in the horizontal axis¹¹⁷.

If the Virgo noise had been perfectly Gaussian and stationary throughout O3, we would expect the rate to decrease exponentially for larger and larger values of the matched-filter SNR, and be independent on the template parameters and particular chunks of data (subsets of the full O3 dataset, each lasting about five days and analyzed as a single block of data). Instead, the rate-vs-SNR curves show a more complicated behavior, with a large variation across the search space and particular data chunks. We observe a non-negligible rate at SNRs as high as 100, while astrophysical signals are typically expected to have SNRs between ~ 5 and ~ 10 . After the application of the χ^2 discriminators, the behavior changes drastically, and the exponential behavior of the rate is recovered, at least as long as we restrict to a subset of the search space. We still observe a large variation of the exponential slope and amplitude across the search space and data chunks, except for the longest templates (green curves of figure 16), which are more robust to instrumental artifacts due to their particular time-frequency signature.

¹¹⁷ At this early stage of the PyCBC analysis, candidates with merger times within fractions of a second from each other can be highly correlated, because a given transient in the data typically ‘rings off’ several templates with high overlaps between each other. The estimated rate of candidates is biased if this correlation is not accounted for. We do so by means of a clustering procedure: a given candidate is ignored if a higher-ranked one exists within a time window of ± 5 s.

The same variation is also seen with candidates from the LIGO detectors, and it is taken into account by the analysis when ranking the multi-detector candidates [58].

Even if properly accounted for, however, a higher rate of noise triggers still reduces our ability to discover compact binaries. Hence, it is informative to inspect the data quality around the triggers in the tail of the curves, in order to understand if a particularly problematic behavior of the detector or analysis can be improved in the future, or if a more robust ranking of the triggers can be found. To this end, we find that the highest SNRs can be attributed to a single segment of ~ 15 min of data on 11 November 2019. These data contain narrow-band, loud and rapidly-varying excesses of power which temporarily affected the data conditioning algorithm used by PyCBC. We determined these excesses to be coming from transient problems with the noise subtraction algorithms used to reconstruct the GW strain channel $h(t)$. Most of the associated high-SNR triggers were automatically removed by the χ^2 discriminators, effectively vetoing the entire problematic segment, however it would have been unlikely for a signal to be detected in Virgo data during this segment. When inspecting the top candidates by χ^2 -weighted SNR, instead, we find that most of the tail is clearly associated with scattered-light glitches.

We conclude from this section that the data conditioning procedure and χ^2 discriminators used by PyCBC, albeit designed for and tuned on LIGO data, are also reasonably effective in Virgo noise. Nevertheless, there seems to be room for further sensitivity gain via a more effective removal of scattered light. One possible way is through better tuning of the Virgo veto stream mechanism introduced in section 4.2, provided that the fraction of affected astrophysical signals can be kept to negligible levels. Another avenue to investigate in the future is a signal-based discriminator or conditioning procedure specifically targeting scattered-light transients, as proposed for example in [59]. Improved tuning and vetoes would have to be evaluated with injection studies and a corresponding calculation of the relative sensitive time-volume of the search.

6.2. Channel safety: channel (in)sensitivity to gravitational waves

Many Virgo data quality analyses aim at ensuring that GW candidates are of astrophysical origin and not caused by terrestrial noises. Typically, searches for correlations between auxiliary channels (monitoring the environment, the detector status, the accuracy of its control, etc) and the $h(t)$ strain channel are run to produce vetoes, that reject times when such correlations are identified. This strategy can lead to a loss of interesting signals if any of the auxiliary channels is sensitive to GWs, which means that it picks up disturbances induced in the detector by these. Hence, a good knowledge of the couplings of auxiliary channels to $h(t)$ is essential. To gather such information, a statistical analysis of all auxiliary channels is performed, using the approach proposed in [60].

This method relies on hardware injections that mimic the effects of GWs on the detector, by moving in a deterministic way one of its test masses. They are used to work around the fact that the transfer functions between $h(t)$ and most auxiliary channels are not well-known, nor understood. The injected signals are 0.6 s long sinusoidal Gaussian functions of various frequencies (between 19 Hz and 811 Hz) and amplitudes (SNR between ~ 20 and ~ 500). The frequencies injected are chosen to scan the entire detection band while avoiding any known resonant frequency (like violin modes). Each waveform is injected three times, spaced by 15 s.

This *safety* analysis assumes that glitches in a given auxiliary channel are distributed according to a stationary Poisson process, whose rate and p -value time series are measured using stretches of data during which no hardware injection is performed. These p -value time series are used to define a classification threshold. Then, a null test is applied to see whether the p -value distribution changes significantly in the presence of hardware injections. Auxiliary

channels that exhibit anomalously small p -values (i.e. lower than the defined threshold) are classified as *unsafe*, meaning that they are likely to mirror excess power coming from the strain channel. The other channels, called *safe*, are the only ones used to produce vetoes.

Virgo DetChar hardware injections were organized at short notice, in the few days between the anticipated end of the O3b run and the moment when the detector had to be switched off because of the pandemic. Among the ~ 2500 auxiliary channels analyzed, 69 were found to be unsafe. That analysis confirmed the existing sets of safe and unsafe Virgo channels, as determined by a previous study. Moreover, its results were in agreement with the safe/unsafe status one could assign *a priori* to auxiliary channels, based on its definitions—that is what quantity they measure, and how these measurements are done. The results validate the Virgo O3 data quality vetoes, both online and offline, which must be based on safe channels to avoid the risk of rejecting real GW signals. The channels identified as unsafe belong to a few well-defined categories: error or correction signals from the DARM control feedback system; correction signals from test mass suspensions; readout channels from the B1 and B1p photodiodes located at the interferometer output (see figure 2); finally, signals monitoring the quality of the detector working point, or used to reconstruct the GW strain channel $h(t)$.

6.3. Spectral noise

The term spectral noise, introduced in section 2.3, identifies the class of detector disturbances appearing as a persistent excess in the noise power spectrum estimation of the data.

Spectral noise has a negative impact especially on searches for persistent GWs, which aim at detecting astrophysical or cosmological signals mainly through the identification of their spectral features. Two typical signal categories of persistent waves are CW [61] and a SGWB [62]. The signals are very weak with respect to the already detected coalescing binary emission. Due to their persistent nature, they can be looked for in the frequency domain where the accumulated power over long observation times can show up at a detectable level, after applying effective signal processing techniques. Moreover, some spectral features of the signals can help in discriminating them from detector noise. On the other hand, spectral noise can mask signals, or produce false candidates, in both cases reducing the search sensitivity.

Searches for persistent signals are typically run off-line, once long stretches of data have been collected. An early identification of spectral disturbances and of their instrumental source would allow to remove, or at least reduce, the source of noise, thus improving the quality of the data.

Different actions can be accomplished at the detector characterization level in support of data analysis. A first action is to identify, and possibly remove, the instrumental source of spectral noises as soon as possible during a data taking period. This is a non trivial task that typically requires a significant amount of work to nail down which detector component is responsible for a given disturbance and to eliminate the noise source, which may imply to replace the noisy component (for instance a cooling fan, an electric motor, etc) [32], to shut it down (if not needed) or to modify it properly. This could consist, for instance, in shifting the frequency of a calibration line which non-linearly couples to another noise source, in order to move the noise line frequency into a band less relevant for GW searches [14, 63].

A second action is the use of additional techniques to differentiate between possible signals and other spectral features. These methods strongly depend on the analysis and on the type of GW signals searched. An example of such techniques relies on the Doppler effect. An

astrophysical CW signal is expected to be modulated in frequency by the Doppler effect, due to the Earth rotation, which induces a shift

$$\Delta f(t) \simeq f_0 \frac{\vec{v}(t) \cdot \hat{n}}{c},$$

where f_0 is the source frequency, \vec{v} the detector velocity, \hat{n} the unit vector identifying the sky direction of the GW source, and c the speed of light. CW searches correct this Doppler effect, thus any monochromatic line present in the $h(t)$ signal is spread by a maximum amount of $\Delta f_{\max} \simeq 10^{-4} f_0 \cdot \cos \beta$, where β is the ecliptic declination. This shift corresponds to up of hundreds or even thousands of frequency bins for typical CW searches.

Potential candidates found in the analysis lead to follow-up investigations to identify a possible instrumental source. This follow-up is also based on a combination of DetChar activity, to spot the source of the disturbances, and application of CW or SQWB algorithms to build confidence in the astrophysical nature of the candidate, see for example [64].

Although spectral noises cannot always be removed, it is still useful to characterize them by constructing a list of noisy lines. This list can be used to exclude those disturbing frequency bands from the analyses, or to veto candidates with frequency too close to those noisy lines

The identification of lines is typically done by automated pipelines [16], based on:

- User-defined thresholds set on data power spectrum or on line *persistence*, defined as the fraction of FFTs, compared to the total number covering the full observation time, in which the ‘normalized’ power content of a given frequency bin was above such a threshold (typically set to six times the average value);
- Highlighting coincidences or significant coherence among different channels;
- Highlighting a pattern in time-frequency maps of the data.

Candidates found in GW searches are subject to verification steps, in which the identification of possible noise counterparts is done by processing the data in the relevant frequency band and period of time and/or running manually one or more of the line identification pipelines. In the following, we report and discuss a few examples of lines identified in Virgo O3 data. Readers can refer to the LIGO-Virgo GWOSC [65] for the official full list of lines.

6.3.1. Combs. Combs are families of lines separated by a constant frequency interval. Typically, noise combs are electromagnetic disturbances generated by digital devices (e.g. microprocessors, programmable communication devices like logical controllers, Ethernet cables, wireless repeaters) that leak into the GW strain channel. Comb lines can have an impact on searches for persistent GWs due to their large number and usually high strength. This makes the identification of combs an important task. There are several combs present in Virgo O3 data, which we describe in the following. They are typically made of many (up to $O(10)$) lines, covering from a few to several Hertz in frequency.

A 1 Hz-spaced comb with 0 Hz offset was already present during previous runs. A new 1 Hz-spaced comb discovered during O3 has a 0.333 Hz offset with respect to integer frequencies. This comb was found during investigations of a line at 22.333 Hz that falls within a region of interest for the Vela pulsar CW search. The instrumental origin of the comb has been confirmed by finding lines at the same frequency in the magnetometers deployed at EGO.

Figure 17 shows the line *persistence* computed over the frequency range 21.8–23.5 Hz on O3 Virgo data. Both 1 Hz combs are clearly visible. Furthermore, there is a comb with




Figure 17. Plot of line *persistence* over the frequency band 20 Hz–30 Hz. The grey box identifies the frequency range covered by a narrow-band search of CW signals from the Vela pulsar in O3a data. The line at 23.333 Hz, which contributed to produce a candidate in the search is clearly visible. In fact, several other lines belonging to the 1 Hz *comb*, both at integer frequencies and shifted by 0.333 Hz, and to a weaker 0.2 Hz-spaced *comb*, are present.

0.2 Hz spacing, whose origin is unknown. The grey shaded area indicates the frequency region explored by a narrow-band CW search targeting the Vela pulsar. The strong line at 22.333 Hz produced an outlier in the search, which was discarded after its instrumental origin was understood.

Finally, two more combs which have been identified by DetChar studies, have both ~ 9.99 Hz spacing, one with 0 Hz offset and the other with 0.5 Hz offset.

6.3.2. Wandering line around 83 Hz–84 Hz. A wandering line is a peculiar kind of spectral noise where the frequency of a spectral line changes with time, with no apparent reason. This is also called a *drifting line* once the mechanism driving the frequency change is at least partially identified, making its variations not entirely random anymore.

An example that triggered lots of DetChar investigations during O3 is the line, normally located between 83 and 84 Hz, as shown in figure 18, that reached about 110 Hz at the maximum of its excursion and had variations of a few Hz over about 1 h [66, 67]. Its origin dates back to the Virgo commissioning run 10 (C10) of August 2018 [68], and possibly even earlier, in the preparatory phase preceding O2 [69]. Neither of the mechanisms that make the line to depart from its typical frequency of about 83 Hz or what produces its variations with time have ever been understood, although several data analysis techniques have been applied and newer ones developed for *line tracking* [68]. An analysis with Bruco [16] revealed no witness channel coherent with $h(t)$ around that line. Moreover, we tracked the frequency evolution of this line, and we correlated the corresponding time series with the auxiliary channels monitoring Virgo [66]. This technique has proven successful in the past, in the case of drifting lines

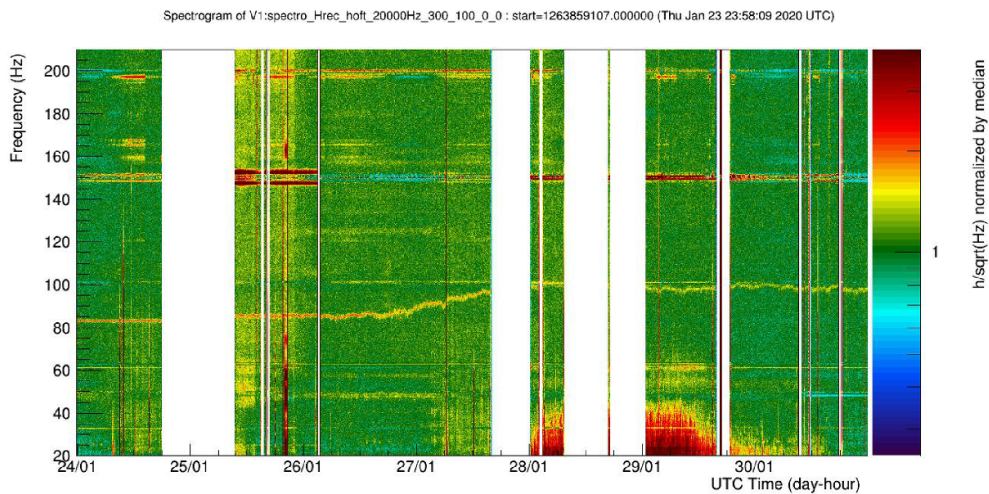


Figure 18. A typical 7 day spectrogram of the $h(t)$ channel, generated by the Spectro tool [16] and allowing the monitoring of a wandering spectral line whose main frequency was first stable around 83 Hz before increasing up to around 100 Hz in about a day.

driven by the temperature of some optical components [70], but has produced no convincing correlation in the case of this line, whose origin remains unknown.

6.3.3. Spectral noise bump around 55 Hz. Figure 19(a) shows the power spectrum of the Virgo GW strain channel $h(t)$ computed at two different dates, 26 February and 2 March 2019 (before the start of the O3 run). Comparing the two plots, one can see that a wide bump around 55 Hz was cured in the meantime. Indeed, a detailed study showed that this disturbance was present most of the time and was observed also in the PRCL channel. This allowed to remove most of this noise excess when producing the reconstructed strain $h(t)$, by accurately subtracting the remaining PRCL contribution [34]. Note that this 55 Hz bump affected the frequencies around 55.6 Hz, where the CW signal possibly emitted by pulsar PSR J1913+1011 is expected. Furthermore, this bump was located within the most sensitive region of the Virgo spectrum for an (isotropic) SQWB search.

6.3.4. Spectral noise around the 50 Hz power line frequency. The GW strain channel $h(t)$ in the frequency region between 45 Hz and 55 Hz was significantly affected by ambient electromagnetic fields originating from the interferometer infrastructure. This noise was studied and mitigated in subsequent steps during the run [32].

The intense 50 Hz line, corresponding to the frequency of the electricity mains, was mitigated and substantially eliminated from $h(t)$ (see figure 19(b)), by implementing a feed-forward noise cancellation scheme using as sensor a voltage monitor of the detector uninterruptible power supply system [32]. This operation did not reduce the 50 Hz harmonics also present in the $h(t)$ spectrum (see figure 8) because they are not due to a non-linear response of the interferometer. They are present in the global environmental disturbances and enter the GW strain channel through different coupling paths.

Sidebands of the mains frequency, at approximately 49.5 Hz and 50.5 Hz, were generated by the pulse width modulation of the electric heater controller of the IMC building. The noise

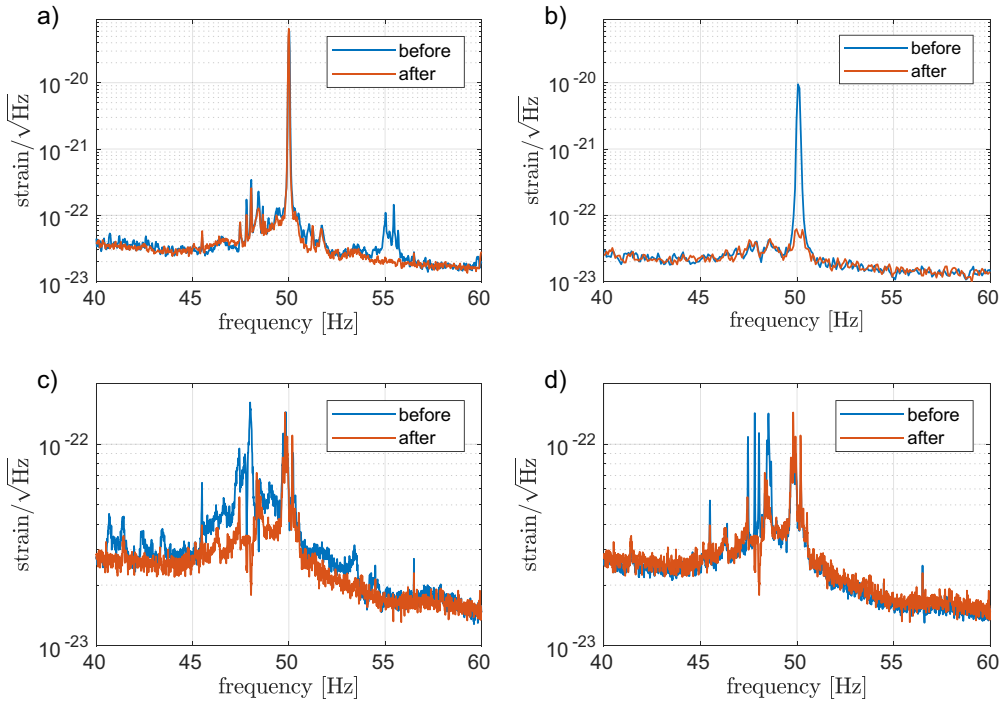


Figure 19. Four steps of the reduction process of the strain spectral noise between 40 Hz and 60 Hz during O3. In all sub-figures the blue curve is before the mitigation, and the red curve is after the mitigation. (a) Cancellation of the spectral noise structure around 55 Hz which was common to the PRCL signal. (b) Subtraction of the 50 Hz line associated to the power grid. (c) Mitigation of a wide-band noise associated to the motor driver crate of the WE suspension. (d) Suppression of a few noise lines associated to mechanical modes of the last stage of the test mass suspension (the payload [41]).

was mitigated by decoupling the electric ground of the building from the central experimental area with an isolation transformer.

Figure 19(c) illustrates a wide-band noise affecting the same region. The origin of this noise was eventually found to be a noisy static voltage accidentally applied to the signal wires of the motors used for positioning and balancing the WE mirror suspension, then coupling capacitively to the mirror coil actuator wires. The noise was mitigated by un-plugging the drivers of the motors, which are not used in Science mode.

Finally, figure 19(d) illustrates a family of lines between 47 Hz and 49 Hz which have been identified as vertical mechanical modes of the last stage of the test mass suspension system. These modes are excited by ambient magnetic fields coupling to the magnetic actuators along the suspension chain. This noise was suppressed by an active mechanical damping of the modes.

6.4. Offline data quality

6.4.1. Offline studies and checks. While section 4.1.2 described the *online* CAT1 vetoes, we focus here on the final set of *offline* CAT1 vetoes. They supersede online vetoes and have been used by all analyses processing the final O3 Virgo dataset. These include analyses using the

Table 7. Virgo O3 offline Science dataset and CAT1 vetoes.

	O3a	O3b	O3a + O3b
Science dataset	12 057 731 s	9611 843 s	21 669 574 s
Logical OR	18 802 s	20 636 s	39 438 s
of all offline CAT1 vetoes	(0.16%)	(0.22%)	(0.18%)

O3 LIGO-Virgo public dataset: that is why the GWOSC website [71] includes detailed public information about these vetoes [72].

Like the online CAT1 vetoes, all these veto flag segments of bad data that are unusable. They are defined with a 1 s granularity and the figure-of-merit used to quantify their impact is their dead time, that is the fraction of Science time that is removed by applying them individually. Yet, the vetoes are not independent and they may overlap. Therefore, they are meant to be applied globally on the dataset, by taking the logical OR of all of them.

Offline vetoes are computed by reanalyzing the whole dataset, using more complex algorithms and additional input information not available online. All the online data quality assessments are reviewed and updated where needed. GPS segments are added to (or removed from) the final dataset by this, possibly iterative, procedure. The offline vetoes defined during the O3 run can be classified into three main categories.

- The duplication—after crosscheck and potential additions or fixes—of online CAT1 vetoes: this includes the saturations of dark fringe photodiodes or mirror suspensions, and the monitoring of the reconstructed GW strain $h(t)$.
- The upgrade of existing online vetoes: the excess rate of glitches is monitored offline using $h(t)$, whereas only the DARM channel could be used online due to latency constraints.
- The addition of new vetoes, based on information that was not available in low latency, or that was not known at the time online flags were generated. These categories are:
 - * Checks of the consistency and of the completeness of the files storing the $h(t)$ GW stream: these vetoes flag segments in which $h(t)$ is missing or contains missing samples.
 - * The $h(t)$ stream is reconstructed by blocks of eight consecutive seconds of data. Therefore, a control loss can possibly impact up to the eight seconds of data that predate it. As the exact time of a control loss is not easy to define, the last ten seconds preceding each recorded control loss have been removed.
 - * The Science dataset has been scanned accurately to identify segments during which the detector was not taking good quality data, contrary to what its status was indicating. These segments were removed from the final dataset.
 - * Finally, a workaround was applied to the detector control system during some weeks in O3b in order to mitigate transient data losses due the failure of an hardware component. That patch allowed to maintain the working point of the instrument, thus sparing a ~ 20 min control acquisition procedure each time it prevented a global control loss. Yet, the application of that workaround could degrade the quality of the data. Thus, the impacted segments were removed from the final dataset, with some safety margin on both ends (the last 10 s before having the control patch be applied automatically, and the first 110 s following the end of the transition back to the nominal control system).

Table 7 summarizes the impact of CAT1 vetoes on the final O3 Science dataset: overall, only 0.2% of the Science data have had to be removed due to various problems.

Conversely, a few minutes of good quality data that had not been included in the online Science dataset for various and clearly understood reasons (software issue, human error, etc) were added to the offline, final, dataset.

6.4.2. Event validation. To assess whether the detection alerts produced by transient searches [49, 50, 73, 74] should be considered as ‘candidate events’, a procedure of *validation* is implemented for each generated trigger [10, 47]. This task has the role to verify if data quality issues, such as instrumental artifacts, environmental disturbances, etc can impact the analysis results and decrease the confidence of a detection, or even foster a rejection [75].

The validation of the online triggers found by GW transient searches includes two separate stages. A prompt evaluation is typically completed within few tens of minutes after an event trigger has been generated, as represented by the data flow in figure 5. Its goal is to determine a preliminary detection confidence and sky localization, in order to deliver public alerts to the astronomy community and support for multi-messenger follow up observations [75], as described in section 5, or to vet that trigger if evidence of severe contamination from non-astrophysical artifacts is present. A team of Virgo DetChar shifters is in charge of this task as part of the RRT (see section 3.1.4). The decision about the event is primarily based on the quick results provided by the DQR, within a few minutes from the trigger. This decision takes into account the evaluation of the operational status of the detector and its subsystems, the environmental conditions, as well as preliminary checks on the strain data. In particular, the shifters are asked to verify the presence of excess noise, namely glitches, around the time of the trigger and the validity of the hypotheses of stationarity and Gaussianity of the data, as discussed in [16]. Moreover, it is examined the possible presence of correlations between the strain data and the auxiliary sensors, which may advise a non-astrophysical origin of the trigger.

With higher latency, a second stage of validation is performed by a *validation team* to finally check candidate events before publications, including those found by offline analyses [51, 76]. Besides of (double-)checking the astrophysical origin of the event trigger, the main purpose of this process is to carefully assess whether the parameter estimation of the source properties can be affected by noise artifacts. This procedure takes advantage of dedicated reruns of the DQR, as well as from additional tools and metrics, including, for example, signal consistency checks [47, 77].

For those events where non-stationary noise, such as glitches, are found in the vicinity of the putative GW signal, or even overlapping with it, a procedure of noise *mitigation* is implemented [78, 79]. During O3b, such process has involved 12 events, including one with Virgo data, GW191105e [10, 80], where the process of mitigation and validation of the data quality has improved the parameter estimation results and credibility. Various O3a events have undertaken a preliminary version of this procedure [8].

7. Preparation of the O4 run

The LIGO-Virgo O3 run has led to the discovery of dozens of new GW signals from compact binary mergers, which have boosted our knowledge of these populations in our local Universe and allowed further, more stringent, tests of general relativity. The O3 run has also been the first long data-taking period for the AdV detector. Thus, it represents a full-scale, extended and non-stop stress test of the organization and work methods of the Virgo DetChar group. The experience accumulated during these 11 months form the base of the DetChar activities, both to prepare and operate for O4 and the following runs.

Although the Virgo DetChar group has fulfilled all its main requirements during the O3 run, work has been going on since then to improve its performance and extend its activities. In particular, the anticipated differences between the O3 and O4 runs lead to new challenges that the group should tackle. The AdV detector will have evolved significantly, with the completion of the Phase I of the AdV+ project [30]. The main changes on the instrument side are the addition of the SR mirror in between the beam splitter and the output port of the Virgo interferometer, a higher input laser power and the implementation of frequency-dependent squeezing. This new configuration will require dedicated instrument characterization activities, while many new data quality features will have to be discovered, understood and later mitigated or solved. On the data analysis side, progress in terms of sensitivity while keeping the network duty cycle high will lead to more GW detections. On the one hand, more work will be required to validate this excess of signal candidates compared to O3. On the other hand, the triggers passing a given false alarm rate threshold will remain dominated by noises, meaning that the bulk of computing resources used by the Virgo DetChar group will not change significantly.

Gathering experience from the past and predictions for the future, a few top priorities have emerged for the DetChar group. A first and obvious one is to broaden the scope of the DetChar monitoring, to make sure that no relevant area remains uncovered, from raw data to the final analyses. Then, the latency of the various DetChar products should be decreased when it is relevant and possible: either by making the corresponding software framework more efficient, or by processing new data more regularly. Finally, some emphasis should be put on increasing the automation of the DetChar analyses and the reporting of their results. In that respect, the DQR is a good example of the realization of these plans. Parallel to common LIGO-Virgo-KAGRA developments on the framework architecture to make DQRs more uniform among the three collaborations and to improve its performance, additional data quality checks will be implemented. They will provide combined results that should give a partial digest of the global vetting of a given GW signal candidate.

The increase of the information available and the help to identify quickly its most relevant points should allow maintaining, if not improving, the high and steady level of Virgo performances observed during O3.

Data availability statement

The data generated and/or analysed during the current study are not publicly available for legal/ethical reasons but are available from the corresponding author on reasonable request. The data that support the findings of this study are available upon reasonable request from the authors.

Acknowledgments

The authors gratefully acknowledge the Italian Istituto Nazionale di Fisica Nucleare (INFN), the French Centre National de la Recherche Scientifique (CNRS) and the Netherlands Organization for Scientific Research (NWO), for the construction and operation of the Virgo detector and the creation and support of the EGO consortium. The authors also gratefully acknowledge research support from these agencies as well as by the Spanish Agencia Estatal de Investigación, the Consellera d'Innovació, Universitats, Ciència i Societat Digital de la Generalitat Valenciana and the CERCA Programme Generalitat de Catalunya, Spain, the National Science Centre of Poland and the European Union—European Regional Development Fund; Foundation for Polish Science (FNP), the Hungarian Scientific Research

Fund (OTKA), the French Lyon Institute of Origins (LIO), the Belgian Fonds de la Recherche Scientifique (FRS-FNRS), Actions de Recherche Concertées (ARC) and Fonds Wetenschappelijk Onderzoek—Vlaanderen (FWO), Belgium, the European Commission. The authors gratefully acknowledge the support of the NSF, STFC, INFN, CNRS and Nikhef for provision of computational resources.

We would like to thank all of the essential workers who put their health at risk during the Covid-19 pandemic, without whom we would not have been able to complete this work.









The authors would also like to thank Samuel Salvador for his extensive and careful proofreading of the manuscript.

ORCID iDs

A Allocca  <https://orcid.org/0000-0002-5288-1351>
A Amato  <https://orcid.org/0000-0001-9557-651X>
N Andres  <https://orcid.org/0000-0002-5360-943X>
M Andrés-Carcasona  <https://orcid.org/0000-0002-8738-1672>
T Andrić  <https://orcid.org/0000-0002-9277-9773>
S Antier  <https://orcid.org/0000-0002-7686-3334>
N Arnaud  <https://orcid.org/0000-0001-6589-8673>
P Astone  <https://orcid.org/0000-0003-4981-4120>
F Aubin  <https://orcid.org/0000-0003-1613-3142>
S Babak  <https://orcid.org/0000-0001-7469-4250>
F Badaracco  <https://orcid.org/0000-0001-8553-7904>
S Bagnasco  <https://orcid.org/0000-0001-6062-6505>
G Baltus  <https://orcid.org/0000-0002-0304-8152>
B Banerjee  <https://orcid.org/0000-0002-8008-2485>
P Barneo  <https://orcid.org/0000-0002-8883-7280>
F Barone  <https://orcid.org/0000-0002-8069-8490>
M Barsuglia  <https://orcid.org/0000-0002-1180-4050>
D Barta  <https://orcid.org/0000-0001-6841-550X>
M Bawaj  <https://orcid.org/0000-0003-3611-3042>
F Beirnaert  <https://orcid.org/0000-0002-4003-7233>
M Bejger  <https://orcid.org/0000-0002-4991-8213>
M Berbel  <https://orcid.org/0000-0001-6345-1798>
S Bernuzzi  <https://orcid.org/0000-0002-2334-0935>
D Bersanetti  <https://orcid.org/0000-0002-7377-415X>
U Bhardwaj  <https://orcid.org/0000-0003-1233-4174>
M-A Bizouard  <https://orcid.org/0000-0002-4618-1674>
R Bonnard  <https://orcid.org/0000-0001-5013-5913>
V Boschi  <https://orcid.org/0000-0001-8665-2293>
V Boudart  <https://orcid.org/0000-0001-9923-4154>
M Branchesi  <https://orcid.org/0000-0003-1643-0526>
M Breschi  <https://orcid.org/0000-0002-3327-3676>
T Briant  <https://orcid.org/0000-0002-6013-1729>
C Buy  <https://orcid.org/0000-0003-2872-8186>
G S Cabourn Davies  <https://orcid.org/0000-0002-4289-3439>
G Cabras  <https://orcid.org/0000-0002-6852-6856>
R Cabrita  <https://orcid.org/0000-0003-0133-1306>

G Cagnoli  <https://orcid.org/0000-0002-7086-6550>
E Capocasa  <https://orcid.org/0000-0003-3762-6958>
G Carullo  <https://orcid.org/0000-0001-9090-1862>
F Cavalier  <https://orcid.org/0000-0002-3658-7240>
R Cavalieri  <https://orcid.org/0000-0001-6064-0569>
G Cella  <https://orcid.org/0000-0002-0752-0338>
E Cesarini  <https://orcid.org/0000-0001-9127-3167>
P Chanial  <https://orcid.org/0000-0003-1753-524X>
E Chassande-Mottin  <https://orcid.org/0000-0003-3768-9908>
S Chaty  <https://orcid.org/0000-0002-5769-8601>
F Chiadini  <https://orcid.org/0000-0002-9339-8622>
A Chincarini  <https://orcid.org/0000-0003-4094-9942>
A Chiummo  <https://orcid.org/0000-0003-2165-2967>
S Choudhary  <https://orcid.org/0000-0003-0949-7298>
N Christensen  <https://orcid.org/0000-0002-6870-4202>
G Ciani  <https://orcid.org/0000-0003-4258-9338>
M Cieřlar  <https://orcid.org/0000-0001-8912-5587>
R Ciolfi  <https://orcid.org/0000-0003-3140-8933>
E Codazzo  <https://orcid.org/0000-0001-7170-8733>
P-F Cohadon  <https://orcid.org/0000-0003-3452-9415>
D E Cohen  <https://orcid.org/0000-0002-0583-9919>
A Colombo  <https://orcid.org/0000-0002-7439-4773>
L Conti  <https://orcid.org/0000-0003-2731-2656>
I Cordero-Carrión  <https://orcid.org/0000-0002-1985-1361>
S Cortese  <https://orcid.org/0000-0002-6504-0973>
M Croquette  <https://orcid.org/0000-0002-8581-5393>
J R Cudell  <https://orcid.org/0000-0002-2003-4238>
T Dal Canton  <https://orcid.org/0000-0001-5078-9044>
S Dall’Osso  <https://orcid.org/0000-0003-4366-8265>
G Dálya  <https://orcid.org/0000-0003-3258-5763>
B D’Angelo  <https://orcid.org/0000-0001-9143-8427>
S Danilishin  <https://orcid.org/0000-0001-7758-7493>
D Davis  <https://orcid.org/0000-0001-5620-6751>
J Degallaix  <https://orcid.org/0000-0002-1019-6911>
S Deléglise  <https://orcid.org/0000-0002-8680-5170>
F De Lillo  <https://orcid.org/0000-0003-4977-0789>
D Dell’Aquila  <https://orcid.org/0000-0001-5895-0664>
A Depasse  <https://orcid.org/0000-0003-1014-8394>
R De Pietri  <https://orcid.org/0000-0003-1556-8304>
R De Rosa  <https://orcid.org/0000-0002-4004-947X>
C Di Giorgio  <https://orcid.org/0000-0003-2127-3991>
F Di Giovanni  <https://orcid.org/0000-0001-8568-9334>
T Di Girolamo  <https://orcid.org/0000-0003-2339-4471>
A Di Lieto  <https://orcid.org/0000-0002-4787-0754>
A Di Michele  <https://orcid.org/0000-0002-0357-2608>
S Di Pace  <https://orcid.org/0000-0001-6759-5676>
I Di Palma  <https://orcid.org/0000-0003-1544-8943>
F Di Renzo  <https://orcid.org/0000-0002-5447-3810>
L D’Onofrio  <https://orcid.org/0000-0001-9546-5959>

M Drago  <https://orcid.org/0000-0002-3738-2431>
D D'Urso  <https://orcid.org/0000-0002-8215-4542>
D Estevez  <https://orcid.org/0000-0002-3021-5964>
F Fabrizi  <https://orcid.org/0000-0002-3809-065X>
V Fafone  <https://orcid.org/0000-0003-1314-1622>
G Favaro  <https://orcid.org/0000-0002-0351-6833>
M Fays  <https://orcid.org/0000-0002-4390-9746>
E Fenyvesi  <https://orcid.org/0000-0003-2777-3719>
I Ferrante  <https://orcid.org/0000-0002-0083-7228>
F Fidicaro  <https://orcid.org/0000-0002-6189-3311>
P Figura  <https://orcid.org/0000-0002-8925-0393>
A Fiori  <https://orcid.org/0000-0003-3174-0688>
I Fiori  <https://orcid.org/0000-0002-0210-516X>
J A Font  <https://orcid.org/0000-0001-6650-2634>
F Frasconi  <https://orcid.org/0000-0003-4204-6587>
A Freise  <https://orcid.org/0000-0001-6586-9901>
G G Fronzé  <https://orcid.org/0000-0003-0966-4279>
B U Gadre  <https://orcid.org/0000-0002-1534-9761>
B Garaventa  <https://orcid.org/0000-0003-2490-404X>
F Garufi  <https://orcid.org/0000-0003-1391-6168>
G Gemme  <https://orcid.org/0000-0002-1127-7406>
A Gennai  <https://orcid.org/0000-0003-0149-2089>
Archisman Ghosh  <https://orcid.org/0000-0003-0423-3533>
B Giacomazzo  <https://orcid.org/0000-0002-6947-4023>
P Giri  <https://orcid.org/0000-0002-4628-2432>
S Gkaitatzis  <https://orcid.org/0000-0001-9420-7499>
B Goncharov  <https://orcid.org/0000-0003-3189-5807>
A Grado  <https://orcid.org/0000-0002-0501-8256>
M Granata  <https://orcid.org/0000-0003-3275-1186>
A Grimaldi  <https://orcid.org/0000-0002-6956-4301>
D Guerra  <https://orcid.org/0000-0003-0029-5390>
G M Guidi  <https://orcid.org/0000-0002-3061-9870>
L Haegel  <https://orcid.org/0000-0002-3680-5519>
O Halim  <https://orcid.org/0000-0003-1326-5481>
J Harms  <https://orcid.org/0000-0002-7332-9806>
A Heidmann  <https://orcid.org/0000-0002-0784-5175>
H Heitmann  <https://orcid.org/0000-0003-0625-5461>
G Hemming  <https://orcid.org/0000-0001-5268-4465>
E Hennes  <https://orcid.org/0000-0002-2246-5496>
V Hui  <https://orcid.org/0000-0002-0233-2346>
B Idzkowski  <https://orcid.org/0000-0001-5869-2714>
P Iosif  <https://orcid.org/0000-0003-1621-7709>
T Jacqmin  <https://orcid.org/0000-0002-0693-4838>
P-E Jacquet  <https://orcid.org/0000-0001-9552-0057>
K Janssens  <https://orcid.org/0000-0001-8760-4429>
P Jaranowski  <https://orcid.org/0000-0001-8085-3414>
C Karathanasis  <https://orcid.org/0000-0002-0642-5507>
S Katsanevas  <https://orcid.org/0000-0003-0324-0758>
S Koley  <https://orcid.org/0000-0002-5793-6665>

M Kolstein  <https://orcid.org/0000-0002-5482-6743>
A Królak  <https://orcid.org/0000-0003-4514-7690>
P Kuijer  <https://orcid.org/0000-0002-6987-2048>
D Laghi  <https://orcid.org/0000-0001-7462-3794>
P Leaci  <https://orcid.org/0000-0002-3997-5046>
M Lenti  <https://orcid.org/0000-0002-2765-3955>
N Leroy  <https://orcid.org/0000-0002-2321-1017>
A Longo  <https://orcid.org/0000-0003-4254-8579>
M Lorenzini  <https://orcid.org/0000-0002-2765-7905>
G Losurdo  <https://orcid.org/0000-0003-0452-746X>
D Lumaca  <https://orcid.org/0000-0002-3628-1591>
C Magazzù  <https://orcid.org/0000-0002-9913-381X>
M Magnozzi  <https://orcid.org/0000-0003-4512-8430>
V Mangano  <https://orcid.org/0000-0001-7902-8505>
M Mantovani  <https://orcid.org/0000-0002-4424-5726>
M Mapelli  <https://orcid.org/0000-0001-8799-2548>
D Marín Pina  <https://orcid.org/0000-0001-6482-1842>
S Marsat  <https://orcid.org/0000-0001-9449-1071>
S Mastrogiovanni  <https://orcid.org/0000-0003-1606-4183>
A Miani  <https://orcid.org/0000-0001-7737-3129>
C Michel  <https://orcid.org/0000-0003-0606-725X>
M Miravet-Tenés  <https://orcid.org/0000-0002-8766-1156>
B Mours  <https://orcid.org/0000-0002-6444-6402>
C M Mow-Lowry  <https://orcid.org/0000-0002-0351-4555>
S Mozzon  <https://orcid.org/0000-0002-8855-2509>
Suvodip Mukherjee  <https://orcid.org/0000-0002-3373-5236>
R Musenich  <https://orcid.org/0000-0002-2168-5462>
I Nardecchia  <https://orcid.org/0000-0001-5558-2595>
C Nguyen  <https://orcid.org/0000-0001-8623-0306>
E Nitoglia  <https://orcid.org/0000-0001-8906-9159>
C Palomba  <https://orcid.org/0000-0002-4450-9883>
F Pannarale  <https://orcid.org/0000-0002-7537-3210>
F Paoletti  <https://orcid.org/0000-0001-8898-1963>
D Pascucci  <https://orcid.org/0000-0003-1907-0175>
R Passaquieti  <https://orcid.org/0000-0003-4753-9428>
B Patricelli  <https://orcid.org/0000-0001-6709-0969>
A Perreca  <https://orcid.org/0000-0002-6269-2490>
K S Phukon  <https://orcid.org/0000-0003-1561-0760>
O J Piccinni  <https://orcid.org/0000-0001-5478-3950>
M Pichot  <https://orcid.org/0000-0002-4439-8968>
L Pierini  <https://orcid.org/0000-0003-0945-2196>
V Pierro  <https://orcid.org/0000-0002-6020-5521>
A Placidi  <https://orcid.org/0000-0001-8032-4416>
W Plastino  <https://orcid.org/0000-0002-5737-6346>
R Poggiani  <https://orcid.org/0000-0002-9968-2464>
E Polini  <https://orcid.org/0000-0003-4059-0765>
R Poulton  <https://orcid.org/0000-0003-2049-520X>
G A Prodi  <https://orcid.org/0000-0001-5256-915X>
M Punturo  <https://orcid.org/0000-0001-8722-4485>

M Razzano  <https://orcid.org/0000-0003-4825-1629>
L Rei  <https://orcid.org/0000-0002-8690-9180>
P Rettugno  <https://orcid.org/0000-0001-8088-3517>
B Revenu  <https://orcid.org/0000-0002-7629-4805>
S Rinaldi  <https://orcid.org/0000-0001-5799-4155>
A Rocchi  <https://orcid.org/0000-0002-1382-9016>
L Rolland  <https://orcid.org/0000-0003-0589-9687>
A Romero  <https://orcid.org/0000-0003-2275-4164>
S Ronchini  <https://orcid.org/0000-0003-0020-687X>
D Rozza  <https://orcid.org/0000-0002-7378-6353>
J Sadiq  <https://orcid.org/0000-0001-5931-3624>
O S Salafia  <https://orcid.org/0000-0003-4924-7322>
F Salemi  <https://orcid.org/0000-0002-9511-3846>
A Samajdar  <https://orcid.org/0000-0002-0857-6018>
N Sanchis-Gual  <https://orcid.org/0000-0001-5375-7494>
A Sanuy  <https://orcid.org/0000-0002-5767-3623>
M Seglar-Arroyo  <https://orcid.org/0000-0001-8654-409X>
Y Setyawati  <https://orcid.org/0000-0003-3718-4491>
N S Shcheblanov  <https://orcid.org/0000-0001-8696-2435>
L Silenzi  <https://orcid.org/0000-0001-7316-3239>
N Singh  <https://orcid.org/0000-0002-1135-3456>
A Singha  <https://orcid.org/0000-0002-9944-5573>
J Soldateschi  <https://orcid.org/0000-0002-5458-5206>
K Soni  <https://orcid.org/0000-0001-8051-7883>
N Sorrentino  <https://orcid.org/0000-0002-1855-5966>
M Spera  <https://orcid.org/0000-0003-0930-6930>
D A Steer  <https://orcid.org/0000-0002-8781-1273>
S Steinlechner  <https://orcid.org/0000-0003-4710-8548>
G Stratta  <https://orcid.org/0000-0003-1055-7980>
A Sur  <https://orcid.org/0000-0001-6635-5080>
B L Swinkels  <https://orcid.org/0000-0002-3066-3601>
E N Tapia San Martín  <https://orcid.org/0000-0002-4817-5606>
A E Tolley  <https://orcid.org/0000-0001-9841-943X>
A Torres-Forné  <https://orcid.org/0000-0001-8709-5118>
I Tosta e Melo  <https://orcid.org/0000-0001-5833-4052>
A Trapananti  <https://orcid.org/0000-0001-7763-5758>
F Travasso  <https://orcid.org/0000-0002-4653-6156>
M Trevor  <https://orcid.org/0000-0002-2728-9508>
M C Tringali  <https://orcid.org/0000-0001-5087-189X>
A Trovato  <https://orcid.org/0000-0002-9714-1904>
K Turbang  <https://orcid.org/0000-0002-9296-8603>
A Utina  <https://orcid.org/0000-0003-2975-9208>
M Valentini  <https://orcid.org/0000-0003-1215-4552>
M van Beuzekom  <https://orcid.org/0000-0002-0500-1286>
J F J van den Brand  <https://orcid.org/0000-0003-4434-5353>
H van Haevermaet  <https://orcid.org/0000-0003-2386-957X>
J V van Heijningen  <https://orcid.org/0000-0002-8391-7513>
N van Remortel  <https://orcid.org/0000-0003-4180-8199>
M Vasúth  <https://orcid.org/0000-0003-4573-8781>

D Verkindt  <https://orcid.org/0000-0003-4344-7227>
A Viceré  <https://orcid.org/0000-0003-0624-6231>
V Villa-Ortega  <https://orcid.org/0000-0001-7983-1963>
M Was  <https://orcid.org/0000-0002-1890-1128>
A R Williamson  <https://orcid.org/0000-0002-7627-8688>
J L Willis  <https://orcid.org/0000-0002-9929-0225>

References

- [1] Abbott B *et al* (LIGO Scientific Collaboration and Virgo Collaboration) 2016 *Phys. Rev. Lett.* **116** 061102
- [2] Aasi J *et al* (LIGO Scientific Collaboration) 2015 *Class. Quantum Grav.* **32** 074001
- [3] Acernese F *et al* (Virgo Collaboration) 2015 *Class. Quantum Grav.* **32** 024001
- [4] Akutsu T *et al* (KAGRA Collaboration) 2021 *Prog. Theor. Exp. Phys.* **2021** 05A102
- [5] Abbott B *et al* (LIGO Scientific Collaboration and Virgo Collaboration) 2017 *Phys. Rev. Lett.* **119** 161101
- [6] Abbott R *et al* (LIGO Scientific Collaboration and Virgo Collaboration) 2021 *Astrophys. J. Lett.* **915** L5
- [7] Abbott B *et al* (LIGO Scientific Collaboration and Virgo Collaboration) 2019 *Phys. Rev. X* **9** 031040
- [8] Abbott R *et al* (LIGO Scientific Collaboration and Virgo Collaboration) 2021 *Phys. Rev. X* **11** 021053
- [9] Abbott R *et al* (LIGO Scientific Collaboration and Virgo Collaboration) 2021 GWTC-2.1: deep extended catalog of compact binary coalescences observed by LIGO and Virgo during the first half of the third observing run (arXiv:2108.01045)
- [10] Abbott R *et al* (LIGO Scientific Collaboration, Virgo Collaboration and KAGRA Collaboration) 2021 GWTC-3: compact binary coalescences observed by LIGO and Virgo during the second part of the third observing run (arXiv:2111.03606)
- [11] LIGO Scientific Collaboration and Virgo Collaboration 2021 Gravitational wave open science center (GWOSC) (available at: www.gw-openscience.org)
- [12] Creighton J D E 1999 *Phys. Rev. D* **60** 022001
- [13] Creighton J D E 1999 *Phys. Rev. D* **60** 021101
- [14] Aasi J *et al* (Virgo Collaboration) 2012 *Class. Quant. Grav.* **29** 155002
- [15] Aasi J *et al* (LIGO Scientific Collaboration and Virgo Collaboration) 2015 *Class. Quantum Grav.* **32** 115012
- [16] Acernese F *et al* 2022 Virgo detector characterization and data quality: tools (arXiv:2210.15634)
- [17] Accadia T *et al* 2012 *J. Instrum.* **7** 3012
- [18] Aasi J *et al* (LIGO Scientific Collaboration and Virgo Collaboration) 2015 *Phys. Rev. D* **91** 022004
- [19] Dooley K L *et al* 2016 *Class. Quantum Grav.* **33** 075009
- [20] LIGO Scientific Collaboration 2015 *Class. Quantum Grav.* **32** 074001
- [21] Abbott B P *et al* 2017 *Phys. Rev. Lett.* **119** 141101
- [22] Abbott B *et al* (LIGO Scientific Collaboration and Virgo Collaboration) 2017 *ApJ Lett.* **848** L13
- [23] Abbott B *et al* (LIGO Scientific Collaboration and Virgo Collaboration) 2017 *ApJ Lett.* **848** L12
- [24] Ballardini G *et al* 2001 *Rev. Sci. Instrum.* **72** 3643–52
- [25] Aisa D *et al* 2016 *Nucl. Instrum. Methods Phys. Res. A* **824** 644–5
- [26] Accadia T *et al* (Virgo Collaboration) 2012 The AdV technical design report *Technical Report VIR-0128A-12* (available at: https://tds.virgo-gw.eu/?call_file=VIR-0128A-12.pdf)
- [27] Travasso F (Virgo Collaboration) 2018 *J. Phys.: Conf. Ser.* **957** 012012
- [28] Acernese F *et al* (Virgo Collaboration) 2019 *Phys. Rev. Lett.* **123** 231108
- [29] Harms J 2019 *Living Rev. Relativ.* **22** 6
- [30] Flaminio R 2021 *Proc. SPIE* **11445** 1144511
- [31] Allocca A *et al* 2020 *Galaxies* **8** 85
- [32] Fiori I *et al* 2020 *Galaxies* **8** 82
- [33] Acernese F *et al* 2022 *Class. Quantum Grav.* **39** 235009
- [34] Acernese F *et al* 2022 *Class. Quantum Grav.* **39** 045006
- [35] Estevez D, Mours B and Pradier T 2021 *Class. Quantum Grav.* **38** 075012

- [36] Estevez D, Lagabbe P, Masserot A, Rolland L, Seglar-Arroyo M and Verkindt D 2021 *Class. Quantum Grav.* **38** 075007
- [37] Graef Rollins J 2016 *Rev. Sci. Instrum.* **87** 094502
- [38] Rollins J G 2014 Advanced LIGO guardian: overview and Coder's introduction *Technical Report* G1400016 (California Institute of Technology) (available at: <https://dcc.ligo.org/LIGO-G1400016/public>)
- [39] Wipf C 2016 SimulinkNb (available at: <https://github.com/cwipf/SimulinkNb>)
- [40] Evans M 2015 Optickle (available at: <https://github.com/Optickle/Optickle/tree/Optickle2>)
- [41] Naticchioni L (Virgo Collaboration) 2018 *J. Phys.: Conf. Ser.* **957** 012002
- [42] Bader M K M 2021 Seismic and Newtonian noise modeling for advanced Virgo and Einstein telescope *PhD Thesis* Vrije Universiteit
- [43] Was M, Gouaty R and Bonnand R 2021 *Class. Quantum Grav.* **38** 075020
- [44] Hild S et al 2009 *Class. Quantum Grav.* **26** 055012
- [45] Fricke T et al 2012 *Class. Quantum Grav.* **29** 065005
- [46] Buikema A et al (aLIGO) 2020 *Phys. Rev. D* **102** 062003
- [47] Davis D et al (LIGO) 2021 *Class. Quantum Grav.* **38** 135014
- [48] DQSEGDB database (available at: www.sciencedirect.com/science/article/pii/S2352711021000224)
- [49] Nitz A H, Dal Canton T, Davis D and Reyes S 2018 *Phys. Rev. D* **98** 024050
- [50] Dal Canton T, Nitz A H, Gadre B, Cabourn Davies G S, Villa-Ortega V, Dent T, Harry I and Xiao L 2021 *Astrophys. J.* **923** 254
- [51] Davies G S, Dent T, Tápai M, Harry I, McIsaac C and Nitz A H 2020 *Phys. Rev. D* **102** 022004
- [52] Davis M H A 1989 *A Review of the Statistical Theory of Signal Detection* (Dordrecht: Springer) pp 73–94
- [53] Abbott B P et al 2020 *Class. Quantum Grav.* **37** 055002
- [54] Allen B 2005 *Phys. Rev. D* **71** 062001
- [55] Nitz A H 2018 *Class. Quantum Grav.* **35** 035016
- [56] LIGO Scientific Collaboration and Virgo Collaboration 2019 *GCN* 26288 (available at: <https://gcn.gsfc.nasa.gov/other/S191124be.gcn3>)
- [57] LIGO Scientific Collaboration and Virgo Collaboration 2020 *GCN* 27306 (available at: <https://gcn.gsfc.nasa.gov/other/S200303ba.gcn3>)
- [58] Nitz A H, Dent T, Dal Canton T, Fairhurst S and Brown D A 2017 *Astrophys. J.* **849** 118
- [59] Tolley A E, Cabourn Davies G S, Harry I W and Lundgren A P 2023 arXiv:2301.10491
- [60] Essick R, Mo G and Katsavounidis E 2021 *Phys. Rev. D* **103** 042003
- [61] Riles K 2017 *Mod. Phys. Lett. A* **32** 1730035
- [62] Christensen N 2019 *Rep. Prog. Phys.* **82** 016903
- [63] Aasi J et al 2015 *Phys. Rev. D* **91** 022004
- [64] Abbott R et al (LIGO Scientific Collaboration, Virgo Collaboration and KAGRA Collaboration) 2022 *Phys. Rev. D* **106** 102008
- [65] LIGO Scientific Collaboration and Virgo Collaboration 2021 O3a instrumental lines (available at: www.gw-openscience.org/O3/o3aspeclines)
- [66] Di Renzo F 2019 Wandering line jumping from 83.5 Hz to 108 Hz (available at: <https://logbook.virgo-gw.eu/virgo/?r=46952>)
- [67] Fiori I, Swinkels B and Di Renzo F 2018 Line at 85 Hz now at 88 Hz and moving (available at: <https://logbook.virgo-gw.eu/virgo/?r=43906>)
- [68] Di Renzo F 2020 Characterisation and mitigation of non-stationary noise in advance gravitational wave detectors *PhD Thesis* Università di Pisa, Pisa, Italy (available at: <https://etd.adm.unipi.it/theses/available/etd-06192020-081853/>)
- [69] Mantovani M et al 2017 84 Hz noise solved (available at: <https://logbook.virgo-gw.eu/virgo/?r=38901>)
- [70] Fiori I et al 2017 Comment to correlation of noise lines drifting in frequency with Env channels (available at: <https://logbook.virgo-gw.eu/virgo/?r=38327>)
- [71] Abbott R et al 2021 *SoftwareX* **13** 100658
- [72] The Virgo DetChar team 2021 Category 1' (CAT1) data quality vetoes applied to the analysis of the O3 run Virgo data (available at: <https://tds.virgo-gw.eu/ql/?c=16808>)
- [73] Adams T, Buskulic D, Germain V, Guidi G M, Marion F, Montani M, Mours B, Piergiorganni F and Wang G 2016 *Class. Quantum Grav.* **33** 175012
- [74] Klimenko S et al 2016 *Phys. Rev. D* **93** 042004

- [75] Abbott B P *et al* 2019 *Astrophys. J.* **875** 161
- [76] Aubin F *et al* 2021 *Class. Quantum Grav.* **38** 095004
- [77] Mozzon S, Nuttall L K, Lundgren A, Dent T, Kumar S, Nitz A H, 2020 *Class. Quantum Grav.* **37** 215014
- [78] Davis D and Walker M 2022 *Galaxies* **10** 12
- [79] Cornish N J and Littenberg T B 2015 *Class. Quantum Grav.* **32** 135012
- [80] LIGO Scientific Collaboration and Virgo Collaboration 2019 *GCN* 26182 (available at: <https://gcn.gsfc.nasa.gov/other/S191105e.gcn3>)