Virgo: a laser interferometer to detect gravitational waves

2012 JINST 7 P03012

(http://iopscience.iop.org/1748-0221/7/03/P03012)

Jinst

PUBLISHED BY IOP PUBLISHING FOR SISSA MEDIALAB

RECEIVED: *December 16, 2011*
ACCEPTED: *February 22, 2012*
PUBLISHED: *March 29, 2012*

Virgo: a laser interferometer to detect gravitational waves

T. Accadia,[11] **F. Acernese,**[5ac] **M. Alshourbagy,**[7ab] **P. Amico,**[6ab] **F. Antonucci,**[8a] **S. Aoudia,**[15a] **N. Arnaud,**[10a] **C. Arnault,**[10a] **K.G. Arun,**[10a] **P. Astone,**[8a] **S. Avino,**[5a] **D. Babusci,**[19] **G. Ballardin,**[2] **F. Barone,**[5ac] **G. Barrand,**[10a] **L. Barsotti,**[7ab] **M. Barsuglia,**[1] **A. Basti,**[7ab] **Th. S. Bauer,**[14a] **F. Beauville,**[11] **M. Bebronne,**[11] **M. Bejger,**[17c] **M.G. Beker,**[14a] **F. Bellachia,**[11] **A. Belletoile,**[11] **J.L. Beney,**[10a,‡] **M. Bernardini,**[7a] **S. Bigotta,**[7ab] **R. Bilhaut,**[10a] **S. Birindelli,**[15a] **M. Bitossi,**[7a] **M.A. Bizouard,**[10a] **M. Blom,**[14a] **C. Boccara,**[10b] **D. Boget,**[11] **F. Bondu,**[15b] **L. Bonelli,**[7ab] **R. Bonnand,**[13] **V. Boschi,**[7a] **L. Bosi,**[6a] **T. Bouedo,**[11] **B. Bouhou,**[1] **A. Bozzi,**[2] **L. Bracci,**[3ac] **S. Braccini,**[7a,†] **C. Bradaschia,**[7a,¶] **M. Branchesi,**[3ab] **T. Briant,**[12] **A. Brillet,**[15a] **V. Brisson,**[10a] **L. Brocco,**[8a] **T. Bulik,**[17bc] **H.J. Bulten,**[14ab] **D. Buskulic,**[11] **C. Buy,**[1] **G. Cagnoli,**[3a] **G. Calamai,**[3ad] **E. Calloni,**[5ab] **E. Campagna,**[3ab] **B. Canuel,**[2] **F. Carbognani,**[2] **L. Carbone,**[6a] **F. Cavalier,**[10a] **R. Cavalieri,**[2] **R. Cecchi,**[7a,§] **G. Cella,**[7a] **E. Cesarini,**[3b] **E. Chassande-Mottin,**[1] **S. Chatterji,**[8a] **R. Chiche,**[10a] **A. Chincarini,**[4] **A. Chiummo,**[2] **N. Christensen,**[2] **A.C. Clapson,**[10] **F. Cleva,**[15a] **E. Coccia,**[9ab] **P.-F. Cohadon,**[12] **C.N. Colacino,**[7ab] **J. Colas,**[2] **A. Colla,**[8ab] **M. Colombini,**[8b] **G. Conforto,**[3ab] **A. Corsi,**[8a] **S. Cortese,**[2] **F. Cottone,**[6ab] **J.-P. Coulon,**[15a] **E. Cuoco,**[2] **S. D'Antonio,**[9a] **G. Daguin,**[11] **A. Dari,**[6ab] **V. Dattilo,**[2] **P.Y. David,**[11] **M. Davier,**[10a] **R. Day,**[2] **G. Debreczeni,**[18] **G. De Carolis,**[7a] **M. Dehamme,**[10a] **R. Del Fabbro,**[7a] **W. Del Pozzo,**[14a] **M. del Prete,**[16b] **L. Derome,**[11] **R. De Rosa,**[5ab] **R. DeSalvo,**[k] **M. Dialinas,**[10a,‡] **L. Di Fiore,**[5a] **A. Di Lieto,**[7ab] **M. Di Paolo Emilio,**[9ac] **A. Di Virgilio,**[7a] **A. Dietz,**[11] **M. Doets,**[14a] **P. Dominici,**[3ab] **A. Dominjon,**[11] **M. Drago,**[16ab] **C. Drezen,**[11] **B. Dujardin,**[15] **B. Dulach,**[19] **C. Eder,**[10a] **A. Eleuteri,**[5a,α] **D. Enard,**[2,†] **M. Evans,**[2] **L. Fabbroni,**[3ac] **V. Fafone,**[9ab] **H. Fang,**[19] **I. Ferrante,**[7ab] **F. Fidecaro,**[7ab] **I. Fiori,**[2] **R. Flaminio,**[13] **D. Forest,**[13] **L. A. Forte,**[5a] **J.-D. Fournier,**[15a] **L. Fournier,**[11] **J. Franc,**[13] **O. Francois,**[2]

[‡]Present address: Subatech 4 Rue Alfred Kastler BP 20722 44307 Nantes CEDEX 3, France.
[†]Deceased.
[¶]Corresponding author; C. Bradaschia.
[§]Present address: Dipartimento di Ingegneria dell'Informazione, Università degli Studi di Siena, Via Roma, 56 53100 Siena, Italy.
[k]Present address: University of Sannio, Benevento, Italy.
[α]Present address: Royal Liverpool Univ. Hosp., Dept. Med. Phys. & Clin. Engn., Liverpool, Merseyside, England.

doi:10.1088/1748-0221/7/03/P03012

2012 JINST 7 P03012

S. Frasca,[8ab] **F. Frasconi,**[7a] **A. Freise,**[2,+] **A. Gaddi,**[7a,*] **M. Galimberti,**[13] **L. Gammaitoni,**[6ab] **P. Ganau,**[13] **C. Garnier,**[10a] **F. Garufi,**[5ab] **M.E. Gspr,**[18] **G. Gemme,**[4] **E. Genin,**[2] **A. Gennai,**[7a] **G. Gennaro,**[7a] **L. Giacobone,**[11] **A. Giazotto,**[7a] **G. Giordano,**[19,β] **L. Giordano,**[5a] **C. Girard,**[11] **R. Gouaty,**[11] **A. Grado,**[5a,††] **M. Granata,**[1] **V. Granata,**[11] **X. Grave,**[11] **C. Greverie,**[15a] **H. Groenstege,**[14a] **G.M. Guidi,**[3ab] **S. Hamdani,**[2] **J.-F. Hayau,**[15b] **S. Hebri,**[2] **A. Heidmann,**[12] **H. Heitmann,**[15] **P. Hello,**[10a] **G. Hemming,**[2] **E. Hennes,**[14a] **R. Hermel,**[11] **P. Heusse,**[10a] **L. Holloway,**[2] **D. Huet,**[2] **M. Iannarelli,**[19] **P. Jaranowski,**[17d] **D. Jehanno,**[10a] **L. Journet,**[11] **S. Karkar,**[11] **T. Ketel,**[14ab] **H. Voet,**[14b] **J. Kovalik,**[6a] **I. Kowalska,**[17b] **S. Kreckelbergh,**[10a] **A. Krolak,**[17ae] **J.C. Lacotte,**[11] **B. Lagrange,**[13] **P. La Penna,**[2] **M. Laval,**[15] **J.C. Le Marec,**[11] **N. Leroy,**[10a] **N. Letendre,**[11] **T.G.F. Li,**[14a] **B. Lieunard,**[11] **N. Liguori,**[16ab] **O. Lodygensky,**[10a] **B. Lopez,**[2] **M. Lorenzini,**[3a] **V. Loriette,**[10b] **G. Losurdo,**[3a] **M. Loupias,**[2] **J.M. Mackowski,**[13] **T. Maiani,**[7a,†] **E. Majorana,**[8a] **C. Magazz,**[7a] **I. Maksimovic,**[10b] **V. Malvezzi,**[9a] **N. Man,**[15a] **S. Mancini,**[7a] **B. Mansoux,**[10a] **M. Mantovani,**[7ac] **F. Marchesoni,**[6a] **F. Marion,**[11] **P. Marin,**[10a,†] **J. Marque,**[2] **F. Martelli,**[3ab] **A. Masserot,**[11] **L. Massonnet,**[11] **G. Matone,**[19,†] **L. Matone,**[10a,μ] **M. Mazzoni,**[3] **F. Menzinger,**[2] **C. Michel,**[13] **L. Milano,**[5ab] **Y. Minenkov,**[9a] **S. Mitra,**[15] **M. Mohan,**[2] **J.-L. Montorio,**[13] **R. Morand,**[11,†] **F. Moreau,**[11] **J. Moreau,**[10b,§§] **N. Morgado,**[13] **A. Morgia,**[9ab] **S. Mosca,**[5ab] **V. Moscatelli,**[8a] **B. Mours,**[11] **P. Mugnier,**[11] **F.-A. Mul,**[14b] **L. Naticchioni,**[8ab] **I. Neri,**[6ab] **F. Nocera,**[2] **E. Pacaud,**[11] **G. Pagliaroli,**[9ac] **A. Pai,**[8a] **L. Palladino,**[9ac] **C. Palomba,**[8a] **F. Paoletti,**[2,7a] **R. Paoletti,**[7a] **A. Paoli,**[2] **S. Pardi,**[5a] **G. Parguez,**[2] **M. Parisi,**[5ab] **A. Pasqualetti,**[2] **R. Passaquieti,**[7ab] **D. Passuello,**[7a] **M. Perciballi,**[8a] **B. Perniola,**[3] **G. Persichetti,**[5ab] **S. Petit,**[11] **M. Pichot,**[15a] **F. Piergiovanni,**[3ab] **M. Pietka,**[17d] **R. Pignard,**[13] **L. Pinard,**[13] **R. Poggiani,**[7ab] **P. Popolizio,**[2] **T. Pradier,**[10a,λ] **M. Prato,**[4] **G.A. Prodi,**[16ab] **M. Punturo,**[6a] **P. Puppo,**[8a] **K. Qipiani,**[5a,π] **O. Rabaste,**[1] **D.S. Rabeling,**[14ab] **I. Rcz,**[18] **F. Raffaelli,**[7a] **P. Rapagnani,**[8ab] **S. Rapisarda,**[7a] **V. Re,**[9ab] **A. Reboux,**[10a] **T. Regimbau,**[15a] **V. Reita,**[10b,δ] **A. Remilleux,**[13] **F. Ricci,**[8ab] **I. Ricciardi,**[5a,φ] **F. Richard,**[2] **M. Ripepe,**[3ac] **F. Robinet,**[10a] **A. Rocchi,**[9a] **L. Rolland,**[11] **R. Romano,**[5ac] **D. Rosiska,**[17cf] **P. Roudier,**[10a] **P. Ruggi,**[2] **G. Russo,**[5a] **L. Salconi,**[2] **V. Sannibale,**[11] **B. Sassolas,**[13] **D. Sentenac,**[2] **S. Solimeno,**[5a] **R. Sottile,**[11]

[+]Present address: School of Physics and Astronomy, University of Birmingham, Edgbaston Birmingham, B15 2TT, U.K.

[*]Present address: CERN, European Organization for Nuclear Research, Geneva, Switzerland.

[β]Present address: Selex Sistemi Integrati, Roma, Italy.

[††]Present address: INAF, Osservatorio Astronomico di Capodimonte, via Moiariello 16, 8013 1 Napoli, Italy.

[μ]Present address: Columbia University 1236 Pupin Hall MC 5247 Box 47 New York NY 10027 U.S.A.

[§§]Present address: Laboratoire Charles Fabry, Institut d'optique, 91127 Palaiseau Cedex, France.

[λ]Present address: IPHC 23 rue du Loess BP 28 67037 Strasbourg CEDEX 2 France.

[π]Present address: E.Andronikashvili Institut of Physics, Tbilisi, Georgia.

[δ]Present address: Institut Neel 25 rue des Martyrs BP 166 38042 Grenoble CEDEX 9 France.

[φ]Present address: CNR–INOA, Istituto Nazionale di Ottica Applicata, Sezione di Napoli, Via Campi Flegrei 34, I-80078 Pozzuoli (NA), Italy.

2012 JINST 7 P03012

L. Sperandio,[9a] **R. Stanga,**[3] **R. Sturani,**[3ab] **B. Swinkels,**[2] **M. Tacca,**[2] **R. Taddei,**[2] **L. Taffarello,**[16c] **M. Tarallo,**[7ab] **S. Tissot,**[11] **A. Toncelli,**[7ab] **M. Tonelli,**[7ab] **O. Torre,**[7ac] **E. Tournefier,**[11] **F. Travasso,**[6ab] **C. Tremola,**[7ab] **E. Turri,**[19] **G. Vajente,**[7ab] **J.F.J. van den Brand,**[14ab] **C. Van Den Broeck,**[14a] **S. van der Putten,**[14a] **M. Vasuth,**[18] **M. Vavoulidis,**[10a] **G. Vedovato,**[16c] **D. Verkindt,**[11] **F. Vetrano,**[3ab] **O. Vziant,**[11] **A. Vicer,**[3ab] **J.-Y. Vinet,**[15a] **S. Vilalte,**[11] **S. Vitale,**[14a] **H. Vocca,**[6a] **R.L. Ward,**[1] **M. Was,**[10a] **K. Yamamoto,**[16bd] **M. Yvert,**[11] **J.-P. Zendri,**[16c] **and Z. Zhang**[2]

[1] *APC, AstroParticule et Cosmologie, Université Paris Diderot, CNRS/IN2P3, CEA/Irfu, Observatoire de Paris, Sorbonne Paris Cité, 10, rue Alice Domon et Léonie Duquet, 75205 Paris Cedex 13, France*

[2] *European Gravitational Observatory (EGO), I-56021 Cascina (PI), Italy*

[3] *INFN, Sezione di Firenze, I-50019 Sesto Fiorentino*[a]*; Università degli Studi di Urbino 'Carlo Bo', I-61029 Urbino*[b]*; Università degli Studi di Firenze, I-50121 Firenze*[c]*, Osservatorio Astrofisico di Arcetri, I-50125 Firenze*[d]*, Italy*

[4] *INFN, Sezione di Genova; I-16146 Genova, Italy*

[5] *INFN, Sezione di Napoli* [a]*; Università di Napoli 'Federico II'*[b] *Complesso Universitario di Monte S. Angelo, I-80126 Napoli; Università di Salerno, Fisciano, I-84084 Salerno*[c]*, Italy*

[6] *INFN, Sezione di Perugia*[a]*; Università di Perugia*[b]*, I-06123 Perugia,Italy*

[7] *INFN, Sezione di Pisa*[a]*; Università di Pisa*[b]*; I-56127 Pisa; Università di Siena, I-53100 Siena*[c]*, Italy*

[8] *INFN, Sezione di Roma*[a]*; Università 'La Sapienza'*[b]*, I-00185 Roma, Italy*

[9] *INFN, Sezione di Roma Tor Vergata*[a]*; Università di Roma Tor Vergata, I-00133 Roma*[b]*; Università dell'Aquila, I-67100 L'Aquila*[c]*, Italy*

[10] *LAL, Université Paris-Sud, IN2P3/CNRS, F-91898 Orsay*[a]*; ESPCI, CNRS, F-75005 Paris*[b]*, France*

[11] *Laboratoire d'Annecy-le-Vieux de Physique des Particules (LAPP), Université de Savoie, CNRS/IN2P3, F-74941 Annecy-Le-Vieux, France*

[12] *Laboratoire Kastler Brossel, ENS, CNRS, UPMC, Université Pierre et Marie Curie, 4 Place Jussieu F-75005 Paris, France*

[13] *Laboratoire des Matériaux Avancés (LMA), IN2P3/CNRS, F-69622 Villeurbanne, Lyon, France*

[14] *Nikhef, Science Park, Amsterdam, the Netherlands*[a]*; VU University Amsterdam, De Boelelaan 1081, 1081 HV Amsterdam, the Netherlands*[b]

[15] *Université Nice-Sophia-Antipolis, CNRS, Observatoire de la Cote d'Azur, F-06304 Nice*[a]*; Institut de Physique de Rennes, CNRS, Université de Rennes 1, 35042 Rennes*[b]*, France*

[16] *INFN, Gruppo Collegato di Trento*[a] *and Università di Trento*[b]*, I-38050 Povo, Trento, Italy; INFN, Sezione di Padova*[c] *and Università di Padova*[d]*, I-35131 Padova, Italy*

[17] *IM-PAN 00-956 Warsaw*[a]*; Astronomical Observatory Warsaw University 00-478 Warsaw*[b]*; CAMK-PAN 00-716 Warsaw*[c]*; Bia lystok University 15-424 Bialystok*[d]*; IPJ 05-400 Świerk-Otwock*[e]*; Institute of Astronomy 65-265 Zielona Gora*[f]*, Poland*

[18] *RMKI, H-1121 Budapest, Konkoly Thege Miklós út 29-33, Hungary*

[19] *INFN, Laboratori nazionali di Frascati, I-00044 Frascati, Italy*

E-mail: carlo.bradaschia@pi.infn.it

ABSTRACT: This paper presents a complete description of Virgo, the French-Italian gravitational wave detector. The detector, built at Cascina, near Pisa (Italy), is a very large Michelson interferometer, with 3 km-long arms.

2012 JINST 7 P03012

In this paper, following a presentation of the physics requirements, leading to the specifications for the construction of the detector, a detailed description of all its different elements is given. These include civil engineering infrastructures, a huge ultra-high vacuum (UHV) chamber (about 6000 cubic metres), all of the optical components, including high quality mirrors and their seismic isolating suspensions, all of the electronics required to control the interferometer and for signal detection. The expected performances of these different elements are given, leading to an overall sensitivity curve as a function of the incoming gravitational wave frequency.

This description represents the detector as built and used in the first data-taking runs. Improvements in different parts have been and continue to be performed, leading to better sensitivities. These will be detailed in a forthcoming paper.

KEYWORDS: Suspensions; Lasers; Interferometry; Antennas

2012 JINST 7 P03012

Contents

1 Introduction **2**

2 Description and requirements **3**
- 2.1 General description 3
- 2.2 Required sensitivity 4
- 2.3 General requirements 5
  - 2.3.1 Seismic isolation 5
  - 2.3.2 Vacuum 5
  - 2.3.3 Interferometer specifications 6
  - 2.3.4 Mirrors 6
  - 2.3.5 Mirror position control 6

3 The infrastructure **7**
- 3.1 Infrastructure function and requirements 7
- 3.2 The site 9
  - 3.2.1 Site choice 9
  - 3.2.2 Site acquisition 10
- 3.3 The central buildings 11
  - 3.3.1 Central Building 11
  - 3.3.2 Mode Cleaner Building and tunnel 13
  - 3.3.3 Control Building 13
  - 3.3.4 Main Building 13
  - 3.3.5 Office Building 14
  - 3.3.6 Technical Buildings 14
- 3.4 The Terminal Buildings 15
- 3.5 The tunnels 16
  - 3.5.1 Arm tunnels 16
  - 3.5.2 Assembly buildings 17
- 3.6 The clean rooms 17
  - 3.6.1 Central Building clean rooms 17
  - 3.6.2 Terminal Building clean rooms 18
- 3.7 Alignment and survey 18
  - 3.7.1 Initial and successive surveys 18
  - 3.7.2 Ground stability 19

4 The vacuum system **22**
- 4.1 Vacuum system function and requirements 22
  - 4.1.1 General requirements 22
  - 4.1.2 Ultimate vacuum 25
- 4.2 The Towers 26

4.2.1 Function and specifications 26
4.2.2 Mechanical design 28
4.2.3 The tower base 29
4.2.4 The Separating Roof 30
4.2.5 The upper cylinder 31
4.2.6 Construction 31
4.2.7 Cleanliness 31
4.2.8 Link tubes 32
4.2.9 Bake-out ovens 32
4.3 The arm tubes 32
4.3.1 Vacuum performances and related constraints 32
4.3.2 Mechanical design of the arm tubes 33
4.3.3 Thermal insulation 39
4.3.4 Bake-out system 40
4.3.5 Tube supports 40
4.4 Mode Cleaner tube 42
4.5 The large valves 44
4.6 Tower and tube assembly 46
4.6.1 Survey 46
4.6.2 Tower positions 46
4.6.3 Tower installation 47
4.6.4 Tube module preparation 48
4.6.5 Tube support installation 48
4.6.6 Tube module positioning and welding 49
4.6.7 Tube leak detection 50
4.6.8 Arm tube geometrical stability 50
4.7 The pumping system 50
4.7.1 Pumping system function and concept 50
4.7.2 Cleanliness 51
4.7.3 Tower pumping 51
4.7.4 Tube pumping 53
4.7.5 Control instruments 54
4.8 Vacuum control system 54
4.8.1 Vacuum control hardware 54
4.8.2 Vacuum control software 56
4.9 Vacuum system performance 57

5 Scattered light noise 59
5.1 Introduction 59
5.2 Noise evaluation 59
5.3 Baffles in the arm tubes 59
5.3.1 Steel baffles 60
5.3.2 Glass baffles 60

5.4 Baffles in the Central Area 61

6 Mirror suspension 64
6.1 Seismic isolation 64
6.2 Superattenuator working principle 64
6.3 The Superattenuator design 66
6.3.1 The Inverted Pendulum top stage 67
6.3.2 The seismic filters 69
6.3.3 The magnetic anti-spring 70
6.3.4 The Filter 7 71
6.4 The mirror payload 72
6.4.1 Mirror suspension wires 73
6.4.2 Payload assembly 75
6.5 The bench payload 76
6.6 Mirror control 76
6.6.1 Mirror position measurement 77
6.6.2 Marionette and Mirror Actuators. 78
6.6.3 Suspension control electronics 79

7 Optics 81
7.1 Introduction 81
7.1.1 Virgo optical scheme 81
7.2 The injection system 82
7.2.1 The 20 W laser 82
7.2.2 External Injection Bench 83
7.2.3 Suspended Injection Bench and Input Mode Cleaner 84
7.2.4 Injection system control 85
7.3 The detection system 87
7.3.1 The detection system design 87
7.3.2 The output mode cleaner 88
7.3.3 The suspended detection bench 88
7.3.4 The external detection bench 89
7.3.5 The end benches 89
7.4 The photodiodes and read-out electronics 90
7.4.1 Photodiodes for the dark fringe and the longitudinal control 90
7.4.2 Photodiodes for the alignment control 90
7.5 7.5 Mirrors 90
7.5.1 Introduction 90
7.5.2 Mirror substrates 92
7.5.3 Mirror coatings 92
7.5.4 The Virgo large mirrors 96

8 Overall interferometer control **99**
8.1 Overall control 99
8.1.1 The hardware environment 100
8.1.2 The timing system 101
8.2 Global interferometer control 101
8.3 The Automatic Locking Procedure 104

9 Environmental monitoring **106**

10 Data acquisition and storage **109**
10.1 Data acquisition (DAQ) 109
10.1.1 Data acquisition choices 109
10.1.2 The front-end read-out 109
10.1.3 The frame builders 110
10.1.4 The frame collectors 110
10.1.5 Raw data flow and latency 111
10.1.6 Data streams 111
10.1.7 DAQ system monitoring 111
10.1.8 DAQ system performance 111
10.2 On-line monitoring 112
10.2.1 Data quality and detector monitoring 112
10.2.2 Data Display 112
10.2.3 Web monitoring 112
10.3 Storage and data management system 112
10.3.1 On-line storage 113
10.3.2 Off-line storage 114
10.3.3 Computing farm 114
10.3.4 Data management 114
10.3.5 Data transfer 114

11 The Virgo effective sensitivity and on-line strain computation **116**
11.1 Detector calibration 116
11.1.1 Mirror actuator transfer function 116
11.1.2 Virgo transfer function and sensitivity 117
11.2 On-Line Preliminary Analysis 117

12 Conclusion **119**

2012 JINST 7 P03012

1 Introduction

Gravitational waves (GW) were predicted by Einstein [1] as a consequence of the theory of General Relativity; its geometrical interpretation states that gravitation is the result of the curvature of space-time caused by mass (or energy) density. In this geometrical view, gravitational waves are thus deformations of space-time which are predicted to travel at the speed of light: by considering small deviations of the metric tensor $g_{\mu\nu}$ from Minkowski metric $\eta_{\mu\nu}$, Einstein's equations can be linearised, giving a wave equation for the quantity $h_{\mu\nu} = g_{\mu\nu} - \eta_{\mu\nu}$ with propagation speed equal to c.

GW emission has been indirectly demonstrated through the radiative energy loss of PSR1913-+16 [2], allowing high precision tests of the predictions of General Relativity. Other discovered binary pulsars successively confirmed the emission of GW. A direct detection would allow for the study of their properties and would initiate a new kind of astronomy; one not based on electromagnetic radiation.

The process of detecting gravitational waves involves measuring directly or indirectly time variations of h; that is variations in the geometry of local space-time. The order of magnitude of the change of distance δL between two free-falling test masses located L meters apart caused by $h_{\mu\nu}$ is

$$\delta L \propto \frac{h}{2} L$$

where h is the typical magnitude of $h_{\mu\nu}$ components.

The main mechanism for generating gravitational waves is acceleration of a mass, in analogy with electromagnetism. Due to the very small value of the gravitational constant G, the amplitude of the generated wave is such that only astrophysical compact objects in violent motion can generate waves with significant amplitude. This amplitude decreases as $1/R$ where R is the distance from the source. Typical GW sources are coalescing binary neutron stars, supernovae and black hole interactions. Nearby events may have a higher amplitude, but are rare, while events occurring far away have lower amplitude but come from a much larger portion of the Universe and are more frequent. Considering events in nearby galaxies at a distance of the order of 20 Mpc, an amplitude of h of the order of 10^{-21} was taken as a reference at the time the interferometer was conceived. For such events the above formula indicates that a precision on the δL measurement of the order of 10^{-18} m for $L = 1$ km is required.

The detection of gravitational waves has been mainly attempted using two approaches. The first involves measuring the amplitude of oscillations of a resonant bar, where the oscillations originate from gravitational waves: following the pioneering work of Weber [3], the Allegro, Auriga, Explorer, Nautilus and Niobe detectors [4] were built and have been operating in a network for several years. The second method measures space-time geometry variations detected by free falling masses using interferometry. Since the 1980s this has seemed to be the most promising technique for detecting gravitational waves. Several detectors based on a Michelson scheme with km-long arms have been built, where the mirrors of the interferometer act as 'free-falling' test masses. Virgo, the French-Italian detector [5], the American LIGO [6], as well as the British-German GEO600 [7] are all now taking data. A smaller interferometer, TAMA, is operating in Japan, where there are also plans for a cryogenic 3 km interferometer (LCGT) [8]. In Australia, the interferometer AIGO [9]

2012 JINST 7 P03012

is being developed. In addition, the LISA project [10] for a space interferometer is also currently being developed.

The Virgo Project was originally proposed by several CNRS and INFN laboratories, which formed the Virgo Collaboration.

Construction and first operation were achieved by CNRS and INFN on the basis of a bilateral agreement. Subsequently, for the running and maintaining of the instrument and the site infrastructures, the European Gravitational Observatory (EGO) Consortium was created by the two funding agencies. In the last few years the Collaboration has been joined by Dutch, Hungarian and Polish groups.

This paper is meant to be a reference for the French-Italian Virgo interferometer, located in Cascina, near Pisa. A comprehensive description of the as-built detector will be presented here, while its detailed performance will be the subject of separate papers.

Section 2 will present the general description and resulting specifications for detector performance. Sections 3, 4 and 5 describe the site and infrastructures and the vacuum chamber, as well as the baffle system that absorbs scattered light. Section 6 describes the suspensions for the seismic isolation of mirrors and optical benches and their position control. Section 7 details the interferometer optical scheme, including the high quality mirrors used in Virgo. Section 8 describes the elements used for interferometer feedback control. The details of the control scheme for longitudinal mirror position and mirror alignment is beyond the scope of this paper. Section 9 describes the system that monitors the environmental conditions that may influence the interferometer. Section 10 describes the data acquisition and on-line data processing systems. Section 11 summarises the expected performance of the Virgo detector, presenting the design sensitivity curve in the form of linear spectral density in the frequency domain. Conclusions and perspectives are presented in section 12.

2 Description and requirements

2.1 General description

The Virgo interferometer is based on the Michelson scheme, which foresees the use of two perpendicular arms of equal length to cancel wavelength fluctuations of the light source. The passage of gravitational waves changes the geometry of space-time, resulting in a different apparent length of each arm. The resulting length difference δL between the two arms gives the gravitational wave signal. In terms of h

$$\delta L = h\, L$$

showing that the response of the instrument to a change h of the metric is proportional to L. The measurement of δL is obtained from the phase difference of the returning beams, which is deduced from the intensity of the interference figure. The best sensitivity is achieved by setting the interferometer at almost complete destructive interference (dark fringe).

Sensitivity is improved by introducing Fabry-Perot resonant cavities in the arms. These multiply the optical path, resulting in a larger phase variation for a given arm length change. However, at high frequency, the gravitational wave may change sign while the photons are stored in the cavity, reducing the phase response of the interferometer. For practical reasons the Virgo arm length was

2012 JINST 7 P03012

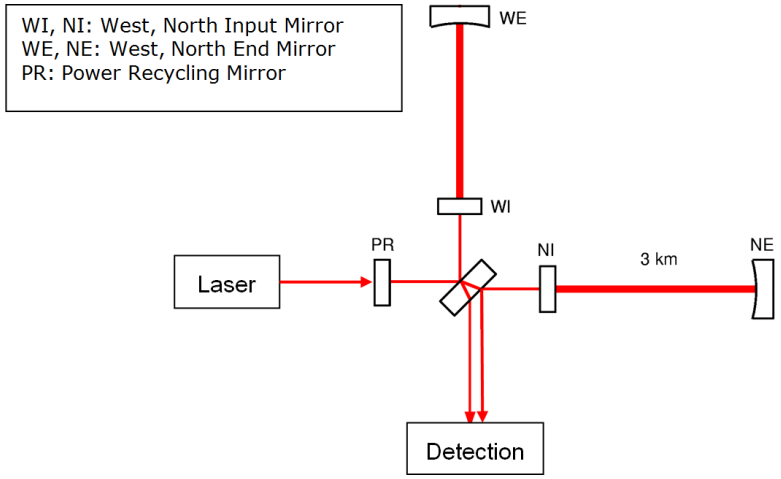

Figure 1. Virgo optical scheme principle.

limited to 3 km, while the optical path increase was chosen to be around a factor of 30, giving an effective optical length of 90 km. The sensitivity of the phase measurement is limited by photon counting statistics and can be improved by increasing the light intensity in the interferometer. By working on the dark fringe, most of the light power goes back to the light source. This light can be usefully sent back to the interferometer, increasing the total circulating intensity by a factor of the order of 50. This is achieved by installing an additional semi-reflective mirror at the entrance of the interferometer that sends back the outgoing light in phase with the incoming beam, forming an overall resonant cavity or 'recycling cavity' [11]. The resulting principle scheme is shown in figure 1.

2.2 Required sensitivity

Due to local gravity fluctuations ground-based interferometers are sensitive to gravitational waves with frequencies higher than a few Hz. At lower frequency a gravitational wave detector has to operate in space. On the high side of the frequency band the gravitational wave signal strongly decreases due to source dynamics. The Virgo interferometer has been designed to cover a frequency band starting from 10 Hz up to a few kHz. Expected sources in this range are coalescing binary systems like neutron star or black hole binaries, stellar collapses, rotating neutron stars, and possibly cosmological background radiation. Estimations of gravitational wave amplitude on the Earth, from typical coalescing neutron star binaries at 20Mpc distance are at the level of $h = 10^{-21}$. This figure was used as a guideline to set a requirement on the spectral sensitivity $\bar{h} \sim 10^{-22}$ to $10^{-23}/\sqrt{\text{Hz}}$.

The Virgo sensitivity is limited by noise sources that can be grouped into different categories. The position of the test masses can fluctuate due to local perturbations, such as residual seismic noise, local gravity fluctuations and thermal motion. Other noise sources affect the detected signal without a real mirror displacement, the main source being photon shot noise. This is illustrated in figure 2, which shows the expected Virgo sensitivity as a function of frequency (in the field of GW detection the sensitivity represents the h spectral amplitude of a wave that would generate the same signal as the noise present in the detector). Typical values are $\bar{h} = 5\ 10^{-23}/\sqrt{\text{Hz}}$ at 100 Hz,

2012 JINST 7 P03012

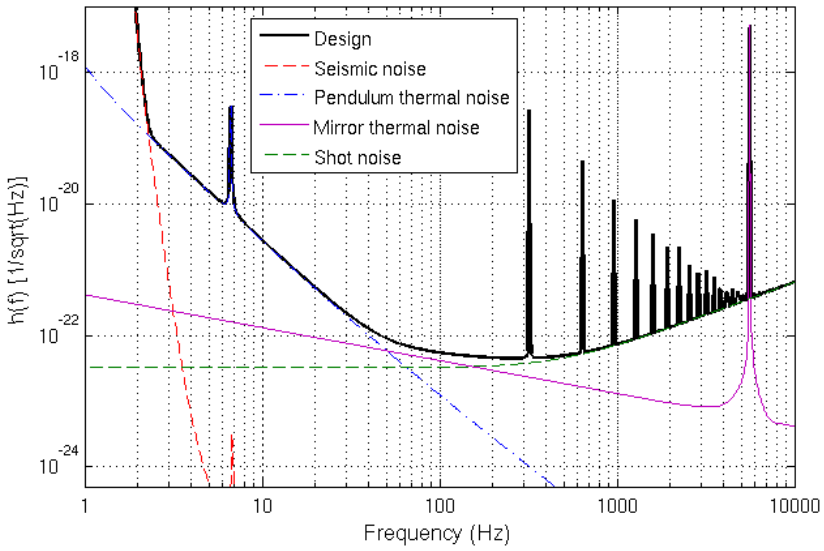

Figure 2. Design sensitivity of Virgo, as a function of frequency. The coloured curves show the main limiting noises.

and $\bar{h} = 3\ 10^{-21}/\sqrt{\text{Hz}}$ at 10 Hz where large efforts have been made to reduce seismic noise through the use of Superattenuators and to thus improve the sensitivity at very low frequency, where many interesting sources can be found.

2.3 General requirements

The construction of a detector of such size is both resource- and time-consuming, particularly in relation to the infrastructure and the vacuum tank. These items were designed to have a life-time of the order of 20 years. It is expected that progress in sensitivity can be achieved with improvements in optics and mirror materials, which can altered rapidly, while using the same infrastructure.

The objective sensitivity induces stringent specifications in the following domains:

2.3.1 Seismic isolation

At the Virgo site seismic noise must be strongly reduced in the detection frequency range (10 to 10000 Hz) up to a factor of $\sim 10^{12}$ at low frequency (10–100 Hz), where it is dominant.

2.3.2 Vacuum

The arm optical length depends on the refraction index along the light beam path, which can change with pressure fluctuations of the medium. To minimise this effect the whole interferometer has to be located in ultra-high vacuum. Moreover, the vacuum isolates the test masses from external acoustic noise, air turbulence and temperature fluctuations. The most stringent requirement comes from the statistical fluctuations of the number of molecules along the beam path. Anticipating a factor of 10 in sensitivity improvement during the detector life-time and keeping another factor of 10 as a safety margin, gives an upper limit for the contribution to the detector noise from this source

2012 JINST 7 P03012

of $10^{-25}/\sqrt{\text{Hz}}$ in $\bar{h}$. With Hydrogen as the main residual gas this translates into a pressure limit of 10^{-7} Pa (10^{-9} mbar).

2.3.3 Interferometer specifications

In a resonant optical cavity, the minimum detectable signal is proportional to $1/Q\sqrt{N}$, Q being the 'quality factor' of the cavity, and N the number of photons entering the cavity, so both values have to be as high as possible. However, the gain introduced by a high value of Q is effective only for signals that have a period longer than the storage time of the cavity. The laser power arriving on the beam splitter mirror of the Michelson interferometer has to be at least 1 kW with a power fluctuation $\delta P/P \leq 10^{-8}/\sqrt{\text{Hz}}$. The laser frequency fluctuations due to coupling to the expected interferometer asymmetries (1%) have to be mastered in order to obtain a laser frequency noise spectral density not larger than 4.10^{-4} Hz/$\sqrt{\text{Hz}}$ at 60 Hz and 3.10^{-4}Hz/$\sqrt{\text{Hz}}$ at 10 Hz.

2.3.4 Mirrors

Each mirror has to meet very stringent specifications: the total losses at the Virgo laser beam wavelength (1064 nm) must be less than 100 ppm. The maximum allowed absorption is 5 ppm to reduce thermal effects, signifying constraints in choice of substrate, and the scattering level has to be less than 5 ppm, requiring very high quality coatings.

The mechanical dissipation of the substrate and coatings determines the level of thermal noise from the mirrors. An overall mechanical quality factor of more than 10^6 is needed.

The wires suspending the mirror are also a source of thermal noise as a result of mechanical dissipation. Also here, low dissipation materials, such as carbon steel or silica, have to be used. A mechanical quality factor for the pendulum of more than 10^6 is required.

2.3.5 Mirror position control

During the whole data-taking operation the interferometer has to remain at its working point, with the two Fabry-Perot and the recycling cavities in a resonant state and the Michelson output on a dark fringe. To reach and keep this situation, suspended optical components need to have their position and angles continuously controlled. The constraints are about 10^{-10} metre rms for the length control (locking), and 10^{-9} rad. rms for the angular degrees of freedom of the mirrors (alignment).

2012 JINST 7 P03012

3 The infrastructure

3.1 Infrastructure function and requirements

The infrastructure for the Virgo interferometer consists of the buildings and the tunnels necessary for the interferometer installation and of all of the technological equipment needed for the operation of the interferometer.

The general requirements used in the creation of the Virgo infrastructure were:

- detector and infrastructure capable of a 20 year life-time;
- detector performances at the limits of present technologies;
- the interferometer to be located on a flat, controlled area, as far away as possible from mechanical vibration sources (such as roads, trains, industries) and within a reasonably short distance of one of the collaborating laboratories;
- perturbation to the geological, biological and economical equilibrium of the surrounding region kept to a minumum.

The following stability requirements were also established:

- the foundations of the buildings have to guarantee that the mirror suspension points will not move more than 1 mm per day;
- the overall displacement of the suspension points in 20 years will be within the design adjustment range of a few centimetres;
- the foundations of the tunnel have to guarantee that the centre of any cross-section of the vacuum tube remains within a straight cylinder of 50 mm radius;
- in order to keep the previous limit, only a limited number of tube supports will have to be realigned, not more frequently than once per year.

The most characteristic Virgo buildings are the long and narrow halls, called tunnels, protecting the 3 km-long vacuum tubes. The two orthogonal tunnels are laid down approximately in the south-north direction and in the east-west direction, starting from an almost cubical Central Building, which contains a large majority of the electronic and optical equipment, including the laser source (figures 3, 4).

At the far ends of the tunnels are the Terminal Buildings, containing the interferometer end mirrors. At the halfway point of each tunnel is a large technical hall, used as the assembly hall for the tubes and, later, as laboratory and workshop space. From the Central Building the Mode Cleaner Tunnel also emerges, ending in the Mode Cleaner Building; these contain the 144 m long Mode Cleaner (MC) resonant cavity.

Close to the Central Building are also:

- the Technical Building, containing the main electricity connections, transformers, uninterrupted power supply, diesel generators, hot and chilled water generators;

2012 JINST 7 P03012

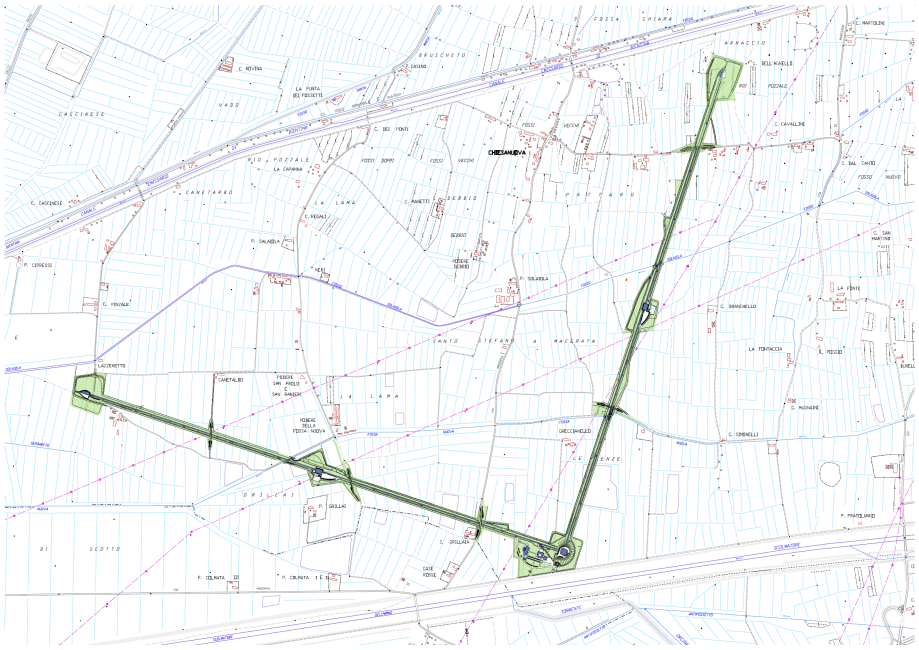

Figure 3. General plan.

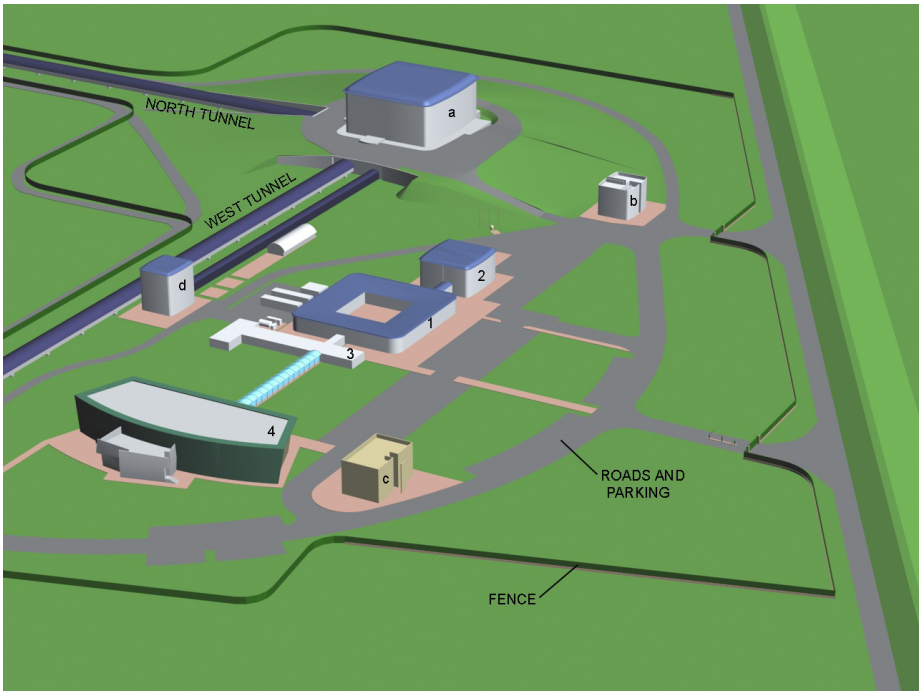

Figure 4. The central site. Building numbering and nomenclature: a- Central Building, b- old Technical Building, c- new Technical Building, d- Mode Cleaner Building, 1- Office Building, 2- Control Building, 3- Office Trailers, 4- Main Building.

2012 JINST 7 P03012

- the Control Building, containing the interferometer Control Room, the main computers and offices;
- the Office Building, added in 2002, containing offices, a seminar room and two smaller meeting rooms;
- the Main Building, completed in 2007, containing offices, laboratories, the auditorium, two meeting rooms and a cafeteria.

Along the arms, on the outside of the 90° angle, is a paved service road, running on the east side of the North tunnel, turning around the central buildings and running on the south side of the West tunnel. The road and the over-crossing bridges are large enough to allow the passage of full-size trucks and buses.

3.2 The site

The land necessary for the interferometer consists of two orthogonal strips, 48 m wide and 3 km long, converging in a central area of about 80,000 m^2 and four areas of about 40,000 m^2 each, at the far ends and at mid-length of the strips. The total surface is about 480,000 m^2. The width of the strips and the dimensions of the five areas were chosen in order to keep external noise sources at a reasonable distance from the interferometer.

3.2.1 Site choice

The Cascina site was preferred for the following reasons:

- the previously listed requirements can be easily met;
- it is flat along the arms, to the level of ± 0.5 m;
- population and farmhouses are of relatively low density;
- there are a few small roads, no highways, no railways.

On this site it has been possible to find an interferometer lay-out that crosses only two secondary roads and keeps Central and Terminal Buildings, containing the mirrors (the test masses), at least 200 m from existing farms and from electrical power lines, and at least 500 m from main roads. Access turned out to be very comfortable through an existing paved road. The whole of the Virgo area is enclosed by a fence, with several access gates; only one of them, close to the central area, is normally operative.

In addition to the two bridges needed for road crossings, three more bridges have been built above the tunnels, in order to allow the circulation of agricultural machines and sheep herds.

The site geology is not ideal from the point of view of stability. In fact the area is an alluvial plane, made of a clay pack several hundred metres deep, and interlaid with a few gravel and sand layers. To prove the suitability of the site for Virgo, a detailed geological survey was performed, including penetrometric tests, core borings and water-table height measurements. Undisturbed core samples were collected down to a depth of 60 m; in some locations 180 m was reached.

2012 JINST 7 P03012

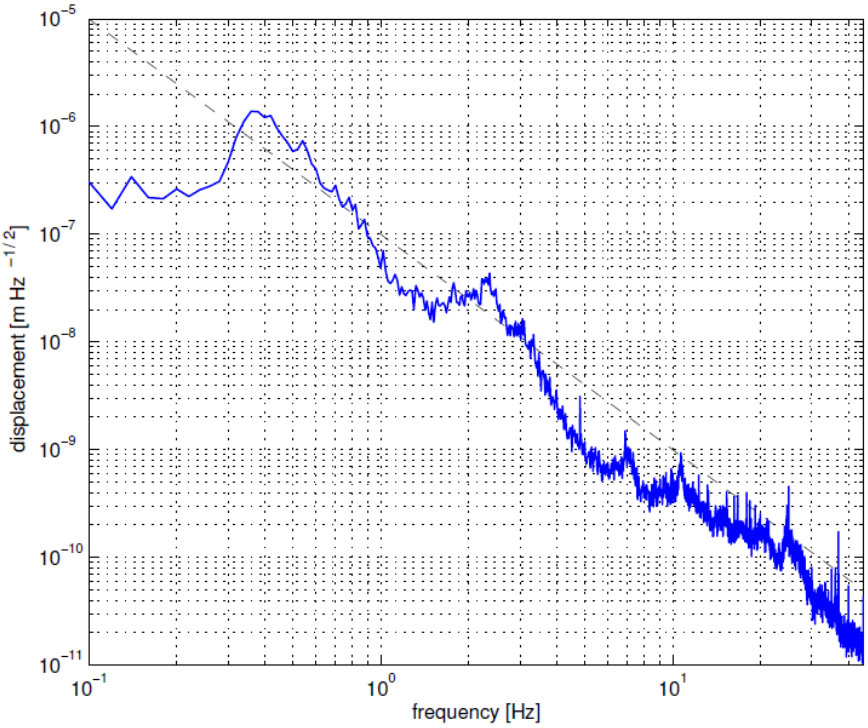

Figure 5. Seismic linear power spectral density.

A computer simulation of the soil subsidence in 20 years was developed by a specialised company, which also used a few measurements taken at a large distance from the site, all available historical data and the existing plans for water pumping.

The results of the study showed an average ground lowering of 15 cm, with a time constant of the order of 5-10 years; the calculated differential subsidence turned out to be below 5 cm, hence within the straightness requirements of the vacuum tubes.

From the seismic point of view, the site turned out to be acceptable, showing a linear power spectrum, as a function of the frequency ν, below the curve $10^{-6}/\nu^2$ m Hz$^{-1/2}$. This condition holds in any weather condition and during peak working hours (figure 5).

In the Cascina region temperature ranges between –5°C, during cold winter nights, and 35°C, on sunny summer days. These data were considered in order to specify the mechanical characteristics of the bellows inserted in the tubes to accommodate thermal expansion.

3.2.2 Site acquisition

Besides the physical characteristics, the Cascina site (via Edoardo Amaldi, Santo Stefano a Macerata, 56021 Cascina, Italia) was chosen thanks to the administrative help of the Mayor and of the County Council of Cascina, of the Province of Pisa, of the Region of Toscany and of Pisa University.

The site was acquired and made ready to build for Virgo by the INFN. The acquisition procedure was a classic 'esproprio', that is a forced acquisition for national interest. Due to the number of owners (about fifty), the operation required a few years to be completed.

2012 JINST 7 P03012

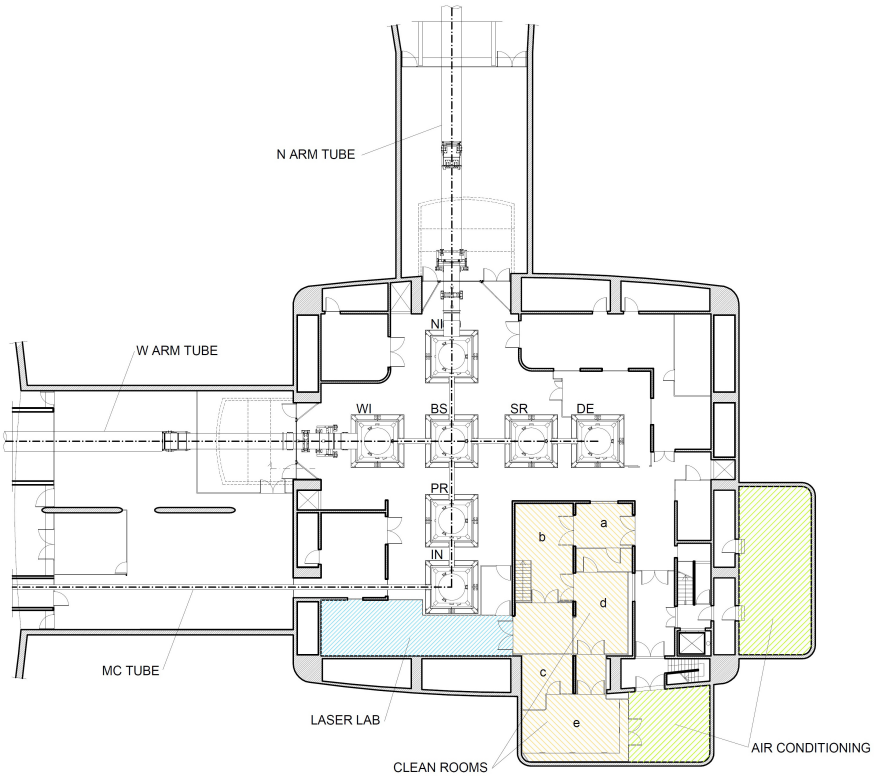

Figure 6. The Central Building ground floor plan. Clean room detail: a- entrance, b- central room, passage to clean basement (trap-door for materials, staircase for personnel), c- class 10 cabin, d- payload assembly, e- washing facilities.

3.3 The central buildings

3.3.1 Central Building

The Central Building (figure 6), at the common starting point of the interferometer arms, contains the main optical components, suspended from six independent Superattenuator (SA) chains, inside six vacuum towers. The bottom part of a seventh tower has been installed, ready for the future installation of a 'Signal Recycling' mirror, along with its SA. It is foreseen that this modification will take place during the transition from Virgo+ to Advanced Virgo (see the Conclusion paragraph for more details). The installation at a later time of a further, 17 tonne-heavy, tower basement, would have been impossible. The seven towers are installed in one single room, with dimensions of approximately 34 m by 34 m, and 18 m high, and is equipped with a bridge crane, which is needed to assemble the SA and the upper parts of the towers. The interferometer beams run 1,100 mm above the floor supporting the towers.

Inside the building, in the south-east corner, there are three floors of laboratories and offices.

On the ground floor, which is the level at which the towers are installed, there is a large area containing clean rooms down to class 10. In one clean laboratory is the laser source; the other clean rooms are used to assemble the optical payloads and for any other operation relating to optical components. On the first floor are: a room for control and acquisition electronics and two

2012 JINST 7 P03012

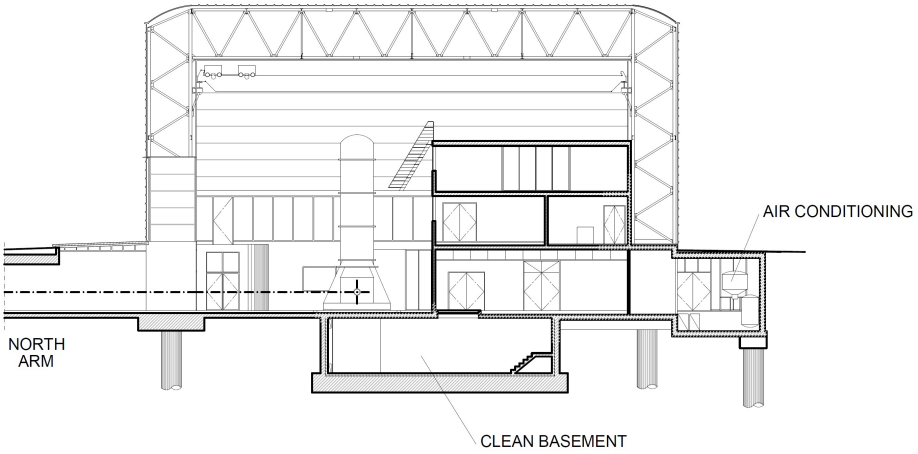

Figure 7. A Central Building elevation cross-section.

electronic laboratories. Optics laboratories and offices are situated on the second floor, while the third floor is open and, being served by the bridge crane, is used as a storage area.

In order to preserve cleanliness, the clean rooms are kept at a slightly higher pressure than the rest of the building, which is, in turn, at higher pressure with respect to the external atmosphere.

On the south and east sides, there are two underground rooms, which are external to the building and which host heat exchangers and air treatment units for the clean rooms and for the general air conditioning system, respectively; hot and chilled water are supplied from the Technical Building.

The Central building (figure 7) has been built with one floor underground, and atop an existing small hill, which is about 4 m high.

In this way the 'ground' floor of the building is at the same level as the floor in the arm tunnels, which emerge from beneath the hill, towards the north and west. The top of the hill provides access to the building at the first floor level, and enables the crossing of the tunnels as if they were under natural bridges. On the South wall are found two doors: one for personnel and one for trucks. Two much larger trap-doors, on the north and west sides, enabled the introduction of the 17 tonne tower basements through the usage of a large mobile crane. Following this operation, those apertures were permanently sealed.

Under the ground floor there is a basement, running beneath all of the towers, and which is accessible through the clean rooms; the basement itself is equipped with filtered air circulation, and has a cleanliness class of between 10,000 and 100,000. The assembled payloads, wrapped in clean foil, are brought from the clean rooms down to the basement, they are then raised into the towers through suitably designed trap-doors; more details on this operation are given in the tower section.

The ground floor and the basement floor have been built as a very rigid double wall concrete box, supported by 25 piles, which run into the clay ground and reach a gravel layer at a depth of about 35 m. In this way the building has very good long-term stability. The upper walls and the roof of the building are made of a 2 m thick light and stiff steel structure, enclosed between two walls of metal sheet and thermal insulation sandwich. This design allowed the pushing of any mechanical resonance of the building above 7.5 Hz, which was necessary as the SA would not have

2012 JINST 7 P03012

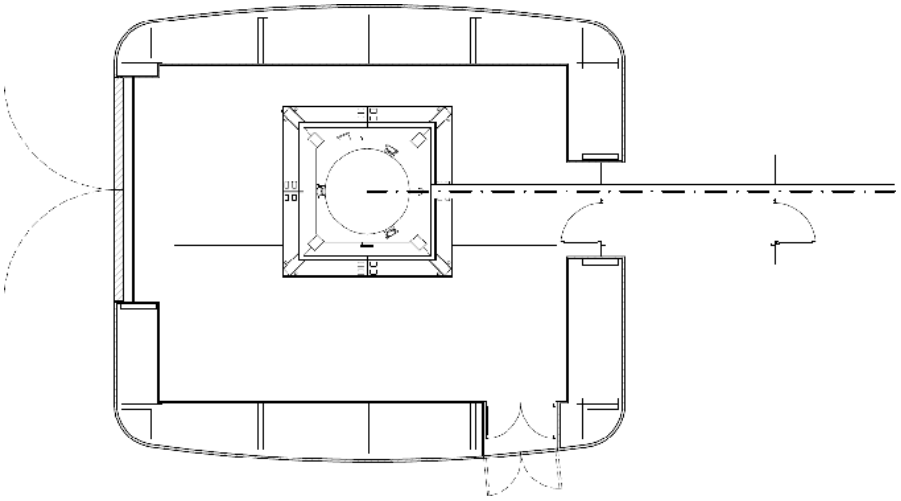

Figure 8. Mode Cleaner Building plan; the MC tower is shown, at the end of the MC vacuum tube.

been able to properly damp vibrations at lower frequencies. The upper layer of the armouring steel embedded in the concrete structure has been completely welded and electrically connected to the upper metal walls, in order to constitute a large Faraday cage, with a good ground connection.

3.3.2 Mode Cleaner Building and tunnel

The MC Building (figure 8) is very similar to the Central Building in terms of structure and main features, but is on a smaller scale, as it only contains one short tower. It has a surface area of 11 m by 12 m and is 12 m high; it has one single room, equipped with a bridge crane, and a clean room basement under the tower, for payload installation.

The MC tunnel, about 130 m long, runs from the Central Building to the MC Building, protecting the MC vacuum tube. The tunnel is 2.7 m wide and 2.7 m high, the lower part is a 'U' shaped concrete beam, supported by 24 m deep pairs of piles, planted every 15 m; the upper part of the walls and the roof are made of metal/insulation/metal sandwich.

3.3.3 Control Building

The ground floor of this building (figure 9) hosts the control room of the interferometer. In this room, displays and projection screens present data and video images coming from the detector, while consoles provide the possibility to fully control the interferometer remotely.

On the upper floor, in addition to some more office space, is the main computing centre for data storage and processing, with about 400 CPU and 250TB of disk. The Network Room hosts the network infrastructure, which covers the whole site, with a total of 180 km of cabling, 130 switches and the internet connection at 1Gb/s.

The building is connected to the Office Building through a covered passage.

3.3.4 Main Building

This is a two floor building, inaugurated at the end of 2007, designed to host personnel and laboratories, which were previously located in prefabricated barracks, more comfortably. It also includes a 99-seat auditorium, two meeting rooms and a cafeteria.

2012 JINST 7 P03012

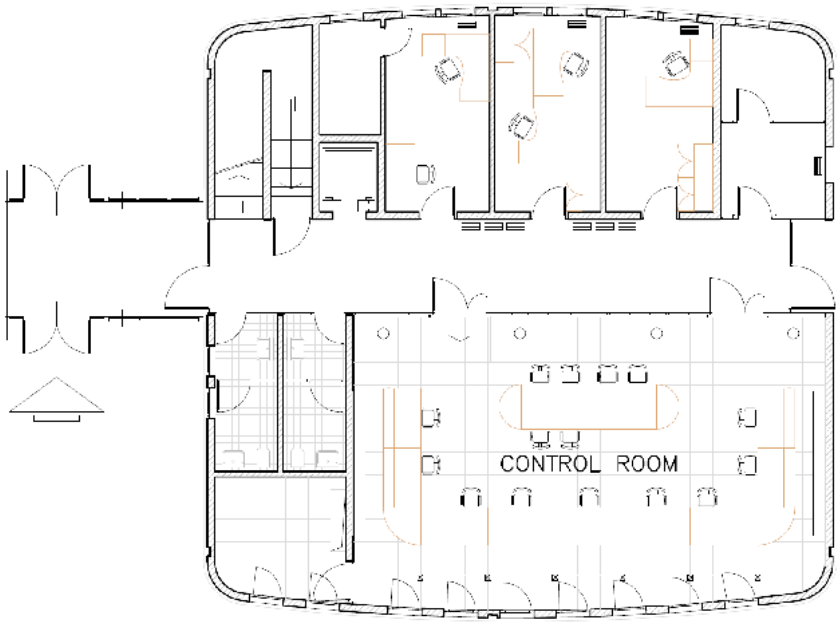

Figure 9. Control Building ground floor plan.

3.3.5 Office Building

The Office Building has the square lay-out of a classical roman villa, with a central '*impluvium*' containing a large olive tree. Within the four sides of the building are two rows of offices, a 50-seat seminar room and two smaller meeting rooms.

3.3.6 Technical Buildings

The main underground connection to the 15 kV, 1 MW power line arrives in the two-floor Technical Building close to the Central Building (figure 4). This building contains all main services for the central zone: transformers, diesel generators, uninterrupted power supplies, hot water generators and, on the roof, chilled water generators.

From the Technical Building, electricity is distributed by underground 15 kV cables to each Assembly Building and to each Terminal Building. At each of the five connection points to the 15 kV line, there are transformers to low voltage (400 V, three-phase). Low voltage available power is: 800 kW in the central zone, 200 kW in each of the four other supply points. From each point power is distributed along the tunnels, up to a distance of 750 m.

Low voltage mono-phase (240 V) and three-phase (400 V) power is distributed to users and plugs at three levels of reliability:

- as from the external company (any amount of interruption time);
- assisted by diesel generators (1 minute maximum interruption time);
- assisted by battery uninterrupted power supplies (0 interruption time).

2012 JINST 7 P03012

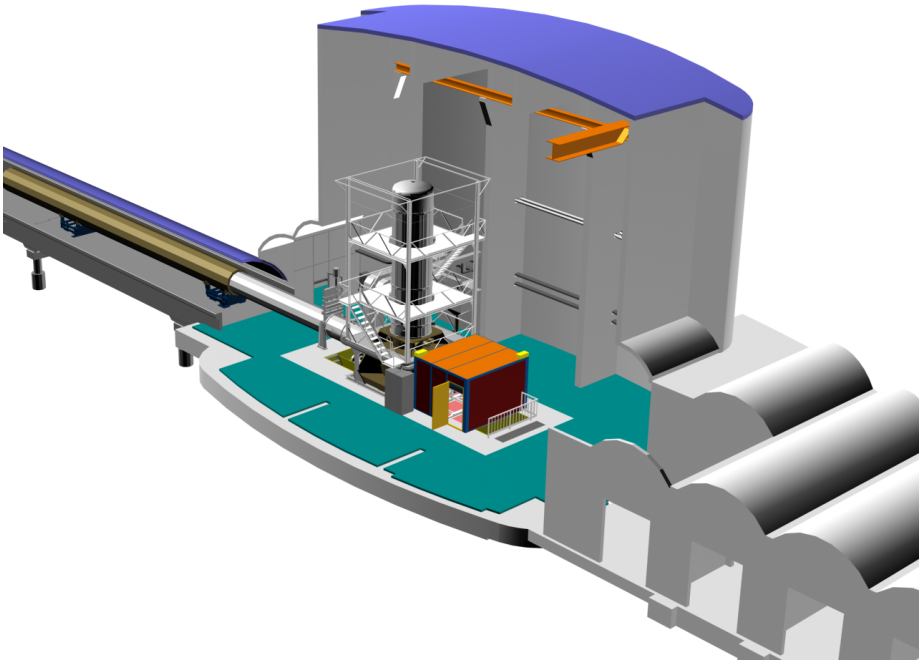

Figure 10. North Terminal Building; the North end tower is shown at the end of the North arm tube. The acoustically isolated cabin behind the tower hides the end optical bench.

Diesel generators can run indefinitely, if supplied with fuel; uninterrupted power supplies can last 30 minutes, but, being assisted by the diesel generators, can last indefinitely.

Respective available powers at the three reliability levels are: 800 kW, 400 kW and 300 kW at the central buildings, 250 kW, 200 kW and 24 kW at the Terminal Buildings and 160 kW, 30 kW and no uninterrupted power, at the mid-arm Assembly Buildings.

A second similar technical building has also been built to serve the new Main building. Its structure allows for the possibility to increase capabilities, in case of potentially larger demands.

3.4 The Terminal Buildings

These are located at the far ends of the interferometer arms and each hosts only one tower, which contains the terminal mirror (figure 10).

All requirements and features are identical to those of the Central and MC Buildings: long-term stability, mechanical resonant frequencies above 7.5 Hz, overpressure for cleanliness, bridge crane, and a clean room basement for payload installation.

However, the requirements have been met with a different design with respect to the Central Building.

Walls and roof constitute a single concrete shell, with strong vertical stiffening ribs on the inside, resting on 30 m deep piles. The covered area is 17 m wide and 25 m long, in the arm direction; the height is 17 m.

The central part of the floor, 6 m wide and 15 m long, supporting the tower and the large valve, is separated by a 5 cm gap from the rest of the floor; this improves the isolation with respect to

2012 JINST 7 P03012

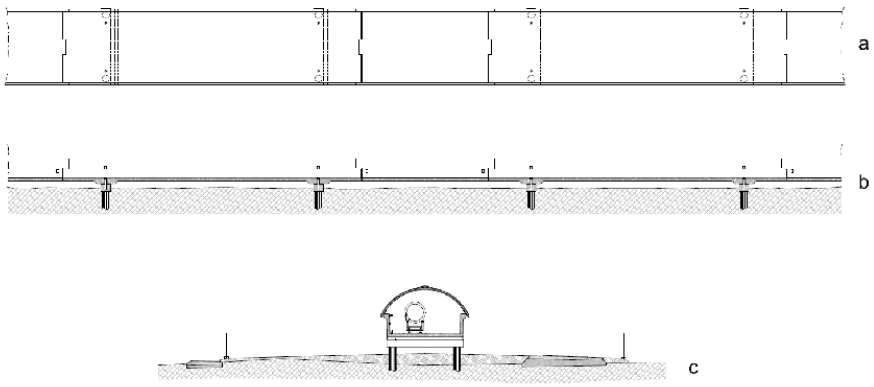

Figure 11. The tunnel; a- concrete beam structure plan, b- concrete beam structure longitudinal cross-section, c- transverse cross-section.

the building vibrations. In order to also improve the stability, this central part of the floor, together with the underlying clean basement, constitutes a very stiff concrete box, founded on a longer set of piles, reaching a more stable gravel layer, 48 m deep at the North arm and 52 m deep at the West arm.

At the back of each building is a one-floor appendix hosting the 15 kV power connection, transformers, uninterrupted power supplies, a diesel generator and chilled/hot water generators for air conditioning.

3.5 The tunnels

3.5.1 Arm tunnels

The function of the tunnels is to protect the arm vacuum tubes from external events of natural and human origin. The most relevant requirement is stability with respect to subsidence, as described at the beginning of the Infrastructure Section. The inner cross-section (height 3.1 m, width 5.0 m) has been designed in order to allow an easy welding of the tube and the passage of a tube module at the side of the already installed tube (figure 11).

The floor and the walls of the tunnel are constituted by concrete beam elements, having a cross-section with the shape of a wide 'U'. The tunnel is covered by a light roof, made of curved sandwich plates (metal/insulation/metal), similar to those used for the central building walls. On the roof there are apertures with filters, to evacuate heat, during tube bake-out.

To ensure a maximum stability, every 15 m a pair of piles, 0.8 m in diameter, has been inserted, 20 m deep, into the ground. On the capitals joining every pile pair, rests the tunnel, having the classical 'Gerber Beam' isostatic structure (figure 11). The concrete beam elements were poured on-site, in a casting factory set up for purpose at the mid-length of the West Arm. There are 20 m long beams resting on two adjacent pile pairs and 10 m long beams joining the longer beams. All the elements come into contact with one another through stiff rubber pads; this helps to avoid mechanical friction between rigid bodies and to damp any possible resonance. Beam resonance effect is, anyway, cancelled out by the fact that tube supports are positioned exactly above the piles, which is the place least affected by oscillations of the structure.

2012 JINST 7 P03012

The tunnel floor has been aligned to better than 1 cm, correcting for Earth curvature, which gives, on a 3 km chord, a sagitta of about 20 cm. The corresponding angle between the vertical directions at the arm ends is about 0.5 mrad.

Main doors to enter the tunnel are located every 300 m, in correspondence with the foreseen positions of the vacuum pumping stations. Further doors, located every 100 m, are present for safety reasons. Between the door wall and the tube there is a 2.5 m wide passage, for equipment and tube modules. On the other side of the tube is a narrower passage, for cable trays and inspection.

3.5.2 Assembly buildings

At the mid-length of each tunnel, there is a large hall about 40 m long and 24 m wide; the whole floor surface, including a 40 m long portion of tunnel, is served by a 5.5 tonne bridge crane. The halls were used for reception, preparation and buffer storage of tube modules, to be directed to the tunnel for installation. Since the end of the tube assembly, the North Assembly building has been converted to a vacuum technology laboratory, a mechanical workshop and storage space. The West Assembly Building became a cryogenic technology laboratory, a silica fibre production laboratory and a monolithic payload assembly laboratory.

3.6 The clean rooms

3.6.1 Central Building clean rooms

On the Central Building ground floor, a surface area of about 300 m^2 is occupied by clean rooms with different purposes (figure 6).

The main clean block is sub-divided into four rooms:

- the entrance/dressing interface, class 10,000;
- the central room (class 10,000) giving access to the other rooms and to the lower clean gallery, under the towers, through a trap-door for equipment and a personnel staircase;
- the assembly laboratory (class 1,000), where payload assemblies can be performed inside an inner volume (class 100), confined by curtains;
- the mirror preparation laboratory, where mirrors are equipped with magnets and markers, inside a class 10 cabin.

On one side-wall of the class 1000 payload assembly room, is the exit door of a large washing/drying machine, which has a 1 m^3 capacity. Components to be cleaned enter the washing machine through the entrance door, accessible from the washing room, which also contains other ultrasonic cleaning facilities and a dedicated ultrapure water production plant.

The clean laser laboratory (class 100,000) has a separate entrance/dressing interface; it contains two optical tables for the laser source and the beam shaping optics. Part of the laboratory walls coincide with the external wall of the Input tower; through the various optical windows, the main beam is injected into the interferometer and control beams reach sensors and cameras.

All the described rooms are supplied by the same air treatment unit, running at full power when there is activity in the clean rooms; during data-taking, in the absence of activity, the clean rooms are set in stand-by, at low power, in order to reduce acoustic noise.

2012 JINST 7 P03012

A dedicated air treatment unit supplies a separate clean room system, constituted by the clean basement and by the towers standing on its roof. The clean basement is the path used to install payloads in the towers. When this operation has to be performed, the bottom lid is removed from the relevant tower, which is supplied at a higher level with filtered air. In this way the arriving payload is subjected to a clean air shower, preventing the entrance of dust into the tower. The air fed to the tower is filtered at a class 100 level, but there is no real laminar flow, due to the uneven air distribution and to the large obstacles in the air path. The air return flow is taken at the basement floor level.

3.6.2 Terminal Building clean rooms

In the Terminal Buildings and in the MC Building, only a clean basement, under the tower, is available for payload installation. This is a small, 20 m^2 clean basement, which is coupled to the tower and functions exactly as the one in the Central Building.

3.7 Alignment and survey

Alignment and survey have been based on a combination of optical instruments and satellite Global Positioning System (GPS). This method allows for the optimisation of measurement time and precision. The used optical instruments are latest generation levels and theodolites, which are fully automated and computer-linked, and able to measure up to distances of the order of several tens of metres. When used alone, such instruments would have incurred rms errors of approximately 15 mm over 3 km, due to the large number of measurement steps necessary to cover the full distance. The introduction, every 300 m, of a GPS measured point enabled the reduction of errors to the level of 5 mm. In this way, successive control surveys in limited regions, do not require the measurement of a full arm, from end to end, but only the measurement of the interval between two adjacent GPS points.

The points to be measured on the building and tunnel floors have been installed with designed-to-purpose precision brass bushes. The bushes have a cylindrical/conical standard hole, into which any kind of target or instrument may be plugged in a reproducible way. About 450 such bushes have been cemented into precision holes or drilled into tunnel and building floors; each bush is closed with a dust-tight protection cover. The tunnel bushes are also considered tube reference points, since the tube is rigidly tied to the tunnel structure, through the supports. Reference points for the towers are determined by four precision holes, drilled at the corners of the 'square plate' of each tower basement.

Precision stands for GPS antennae have been developed in order to transfer the position of floor bushes to outside of the tunnel, through suitable holes in the tunnel roof.

3.7.1 Initial and successive surveys

A first survey of the site was performed with the help of experts from the Engineering faculties of Roma-1 and Pisa universities, referring to IGM (*Istituto Geografico Militare* – Military Geographical Institute) points over a 50 km-wide region. Several secondary reference monuments were built on the site, to be used during tunnel and building construction.

The main tunnel surveys were performed on the still open concrete structure, before they were covered with roof panels. All the reference bushes, located every 15 m, in correspondence with

2012 JINST 7 P03012

the tube supports, were carefully surveyed (± 2 mm, rms). Subsequently, tube supports and tube modules were positioned with respect to the bushes (± 1 mm, rms).

Since the tube installation, all of the bushes have been surveyed at least once a year, to monitor tunnel/tube (ground) movements. When the displacements are larger than given limits (see the Vacuum System Section) opposite correction displacements are performed, using the adjustment capabilities of the tube supports.

3.7.2 Ground stability

Ground lowering rate at the Virgo site turned out to be of the expected order of magnitude: between 0 and a maximum of 30 mm per year (figure 12).

In reality, only a limited saturation effect can be seen, even about seven years after construction (figure 13).

In general the largest subsidence effects are located where the load added on the ground is large, i.e. in correspondence with the Terminal Buildings, the mid-arm assembly buildings, the bridges and earth fills.

The Central Building, probably thanks to its deep foundations reaching a consistent gravel layer, has remained stable, together with a large fraction of the first half of the North tunnel.

The Terminal Buildings, despite similar foundations on deep gravel, sink at a relatively fast rate: 8 mm per year at the North End and 12 mm per year at the West. The end slabs of the tunnels, having less well-performing foundations, sink twice as quickly as the respective Terminal Buildings (figure 14).

This behaviour produces a vertical share strain on the last tube module, rigidly connected to the large valve. In order to keep this strain within acceptable limits, the last tube modules of each arm have to be realigned every few months, this effectively means every 2 mm of displacement of the tunnel-end with respect to the Terminal Building. At the tunnel/Terminal Building interface, a multi-point hydraulic level system has been installed, to continuously monitor the relative displacements at 0.1 mm level.

Horizontal transversal ground displacements are, as expected, more than 10 times smaller than the vertical ones. Nevertheless, in a few points it has been necessary to also perform horizontal tube adjustments.

Horizontal longitudinal displacements are also expected to be small and neutralised by the interaction between tunnel beam elements.

Even if the lowering rate does not decrease, the adjustment capabilities of the tube supports will allow for Virgo operation for more than 20 years.

2012 JINST 7 P03012

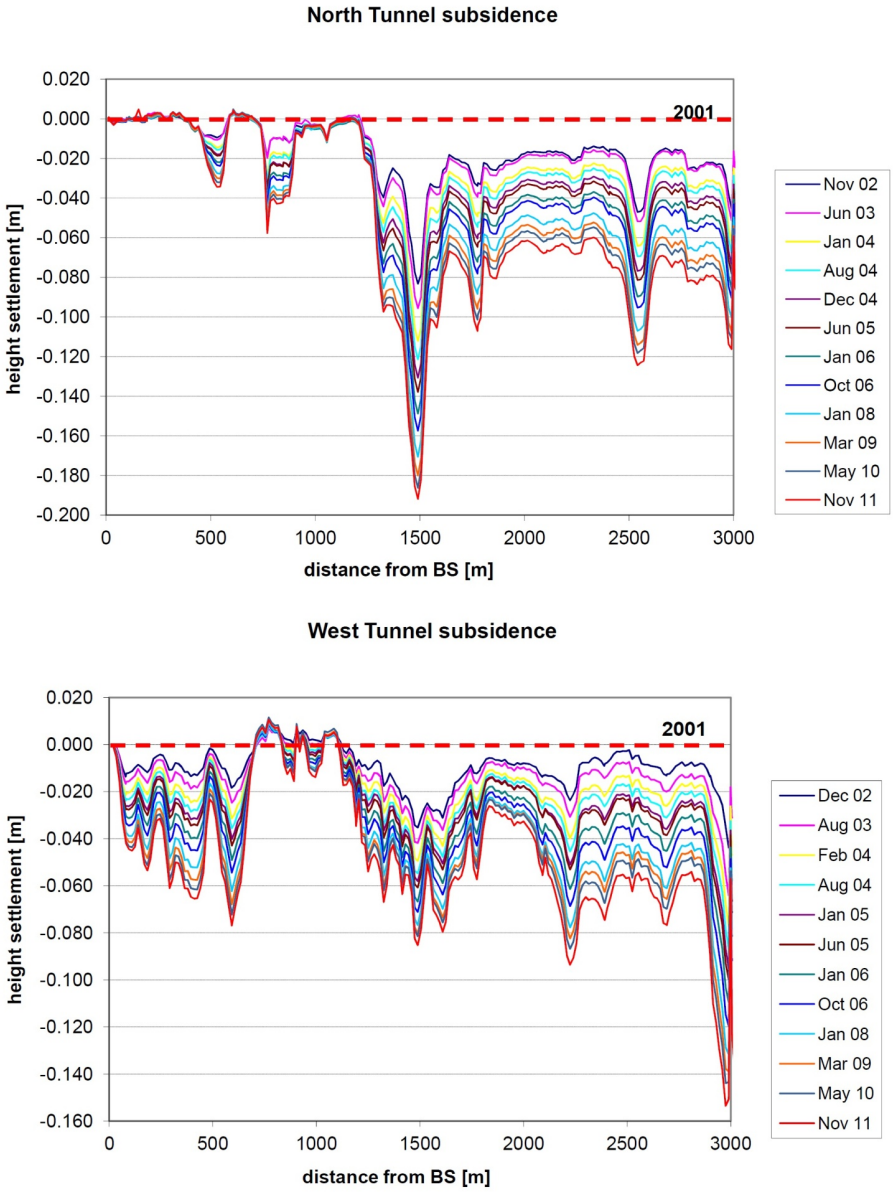

Figure 12. Subsidence profiles of the North and West Arms.

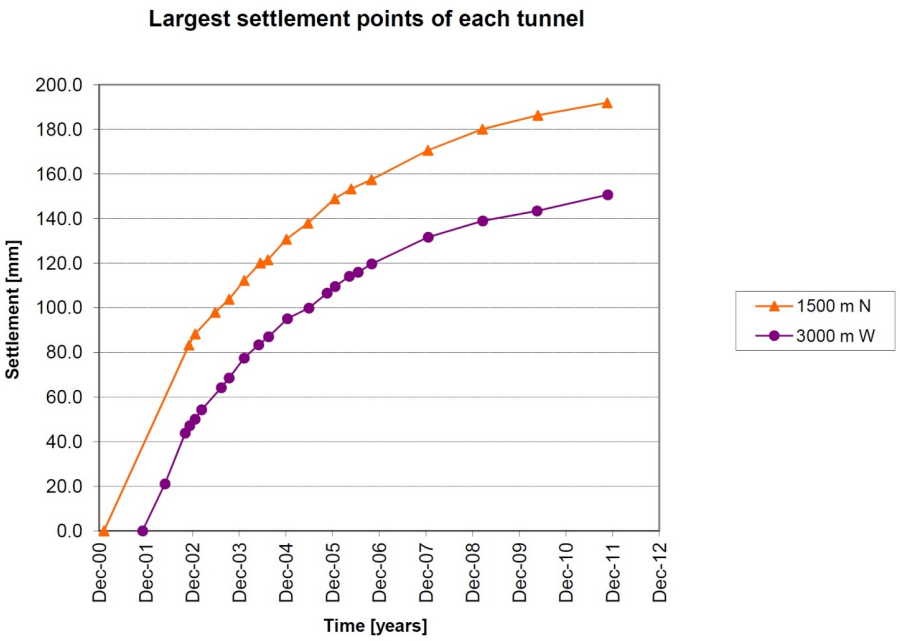

Figure 13. Subsidence time evolution.

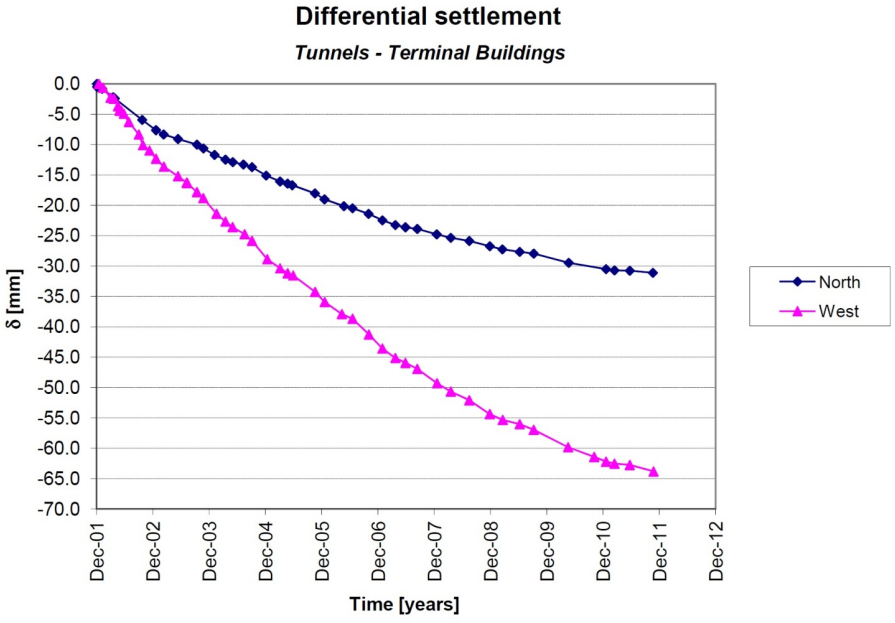

Figure 14. Tunnel-sinking relative to the Terminal Buildings.

4 The vacuum system

4.1 Vacuum system function and requirements

4.1.1 General requirements

The entire interferometer is inside a large vacuum enclosure, in order to suppress several sources of noise, each one of them being strong enough, at atmospheric pressure, to prevent the observation of gravitational waves.

The vacuum vessel consists of several parts:

- two 3 km-long straight tubes, 1.2 m in diameter, lying on the ground and perpendicular to one another; the laser beams propagate along the interferometer arms inside these tubes;
- several vertical cylinders, known as 'towers', 2 m in diameter and up to 11 m high, containing the optical elements suspended to anti-seismic attenuators - the Superattenuators;
- one 142 m-long tube, 30 cm in diameter, containing the Mode Cleaner optical cavity;
- several valves, with apertures of up to 1 m;
- several pumping groups, including many types of pumps and vacuum gauges.

The chosen diameter of the 3 km-long arm tubes was 1.2 m, in order to be able to contain two or even three independent interferometer beams. In this way the implementation of an additional interferometer would require major changes only to the towers.

The 1 m aperture gate valves at both ends of each 3 km tube, known as the 'Large Valves', were developed based on a custom design. Their fundamental role consists in separating the towers from the tubes, enabling the tubes to be kept in vacuum, while venting of the towers takes place; in this way it is possible to save a large amount of work and money, allowing frequent maintenance operations in the towers.

The noise sources suppressed by vacuum are:

- transmission of acoustical noise to the mirrors;
- thermal noise due to damping of mirror suspensions by air friction;
- excitation of mirror motion by gas molecules;
- scattering of light by residual gas molecules;
- and refractive index changes due to statistical fluctuations of gas molecule density.

It is apparent that the last phenomenon is dominant. It is therefore sufficient to reach a pressure low enough to suppress this noise and thus ensure that all of the other ones are eliminated as well.

The noise level superimposed to the gravitational wave signal by residual gas pressure fluctuations has been carefully evaluated. The maximum residual pressure acceptable value has been chosen in order to keep this noise lower by one order of magnitude with respect to the shot noise. This condition must hold also for a future upgraded interferometer, having a sensitivity improved by a factor of 10, with respect to the Virgo design sensitivity.

2012 JINST 7 P03012

Table 1. Average partial pressures on the optical path of the interferometer beams.

Gas species	Partial pressure
H_2	10^{-7} Pa (10^{-9} mbar)
all other gases	10^{-8} Pa (10^{-10} mbar)
hydrocarbons	10^{-11} Pa (10^{-13} mbar)

Quantitatively, assuming a design sensitivity of 10^{-23} $Hz^{-1/2}$ above 300 Hz, and 10^{-24} $Hz^{-1/2}$ for an improved detector, the noise signal due to pressure fluctuations must be below 10^{-25} $Hz^{-1/2}$. This corresponds to a residual gas pressure of about 10^{-7} Pa (10^{-9} mbar). This value has been calculated for a residual gas constituted essentially of Hydrogen, which is the case for baked stainless steel vacuum enclosures. In the case of residual gas dominated by higher polarisability gases, the limit pressure should be lowered by one order of magnitude.

Baked vacuum systems were considered, as, in order to reach the requested end pressure, it is mandatory to foresee a bake-out process for the whole tank, when it is already evacuated. Experimenting with prototypes allowed us to choose the most appropriate bake-out parameters to eliminate the H_2O molecular layers present on the inside surface of tubes and towers. The best choice turned out to be heating the tanks, once under vacuum, to 150°C for about 100 hours.

Furthermore, the residual gas must be free of condensable organic molecules (hydrocarbons), in order to keep the optical surfaces clean. A hydrocarbon partial pressure of 10^{-11} Pa (10^{-13} mbar) is required should one want to avoid the cumulative deposition of one monolayer of molecules on the optical elements over a 4 year period. This value was chosen theoretically, because of the lack of experimental data. A reasonable sticking probability of 0.1 for hydrocarbon molecules was used, together with other more conservative hypotheses:

- that one monolayer of molecules deteriorates mirror surfaces;
- that hydrocarbon molecules polymerise under the intense Fabry-Perot beams (10 kW), while it is more likely that they are photo-desorbed.

The retained target values for the average partial pressures on the optical path of the interferometer beams are listed in table 1.

These very stringent conditions (ultra-high vacuum) are necessary on the 3 km optical paths and in the lower part of the towers containing the main optical components. A 10^3 times higher pressure (high-vacuum (HV)) is acceptable in the towers containing the Input and Detection optical benches, the Mode Cleaner cavity and in the upper part of all the towers, where most of the electromechanical parts of the SA are contained. In order to cope with a HV regime (10^{-4} Pa), a careful material selection has been performed for all SA components, and suitable vacuum/thermal treatments have been developed. UHV volumes and HV volumes are separated by vacuum-tight optical windows or by low-conductance apertures, allowing the preservation of the pressure difference when operating in a differential vacuum regime.

Pressure fluctuations in the vacuum enclosure can also be produced by bursts of molecules released by the vessel walls or by internal components. The origin of bursts can be mechanical (deformation or friction) or thermo-electrical (discharges in ion pumps). Molecule burst releases

2012 JINST 7 P03012

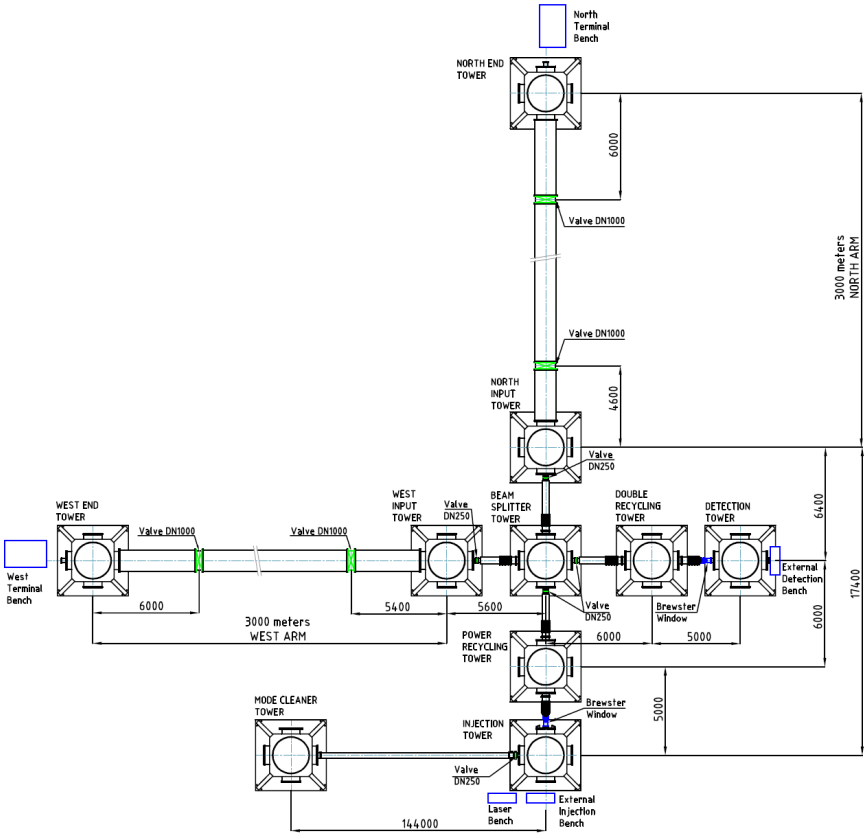

Figure 15. Schematic lay-out of the vacuum system.

are reduced to the very minimum by the absence of moving parts in the UHV volumes, pumped only by Titanium evaporation pumps and Ion pumps, and by the extreme cleanliness of components, allowing a very low load on the pumps themselves.

Residual gas burst effects, still present in the gravitational wave signal, will be vetoed by their peculiar time structure. Their shape should be characterised by a rise time well below 1 s [12], due to the fast emission process, and by a decay time larger than 100 s, due to the ratio V/S of the volume V over the pumping speed S of the corresponding pumping group.

Figure 15 shows a schematic lay-out of the vacuum system.

Finally, the evacuated volume has to be kept free of dust, in order to avoid optics contamination. This request holds, in particular, for the tower lower part, where optical components are located. Quantitatively, the fraction of mirror surface covered by dust grains should be kept at the ppm level, which is the same level prescribed for scattered light losses.

In order to cope with all the requirements, it was decided that the material for tower and tube construction must be stainless steel 304L. Possible alternatives could have been: higher quality stainless steels or aluminum alloys. 304L stainless steel was preferred, after preliminary dimensioning of the various alternatives, as being the least expensive and most reliable solution. The following main points were considered:

2012 JINST 7 P03012

- mechanical characteristics;
- cost and availability of the raw material;
- welding reliability and crack-less deformability;
- raw material outgassing rate and possible reduction processes;
- construction experience within the CNRS and INFN;
- experience inside candidate construction companies;
- compatibility with existing vacuum instrumentation.

Among the different stainless steel grades, 304L was considered the best compromise between cost, ease of procurement and absence of ferromagnetic behaviour.

The influence of diffused light on the interferometer depends strongly on vacuum enclosure properties. In fact, stray light can be diffused back into the interferometer beams after bouncing off vacuum enclosure walls; in this way seismic modulation can spill into the gravitational wave signal. To avoid these problems a thorough investigation of diffused light propagation has been performed, as described in the relevant section. As a result, a large set of baffles and absorbers has been installed inside the vacuum volume, with the aim of cutting the path of diffused light from mirrors to vacuum walls and back to the mirrors. Stainless steel baffles have been placed inside the 3 km arm tubes. Inside the tower volume, up to the large valve limit, there are also baffles made of special infrared absorbing glass.

4.1.2 Ultimate vacuum

As explained in the introductory section, an end pressure below 10^{-7} Pa has to be reached on the light beam path.

This pressure value becomes a real challenge if it has to be reached in a large volume, with kilometric dimensions and a huge wall surface. In a vacuum enclosure without inner gas sources, with negligible permeation effects and without leaks, the end pressure close to the pumping ports is determined by the ratio:

$$\mathrm{end}_{\mathrm{pressure}} = \frac{\mathrm{wall}_{\mathrm{surface}} \times \mathrm{outgassing}_{\mathrm{rate}}}{\mathrm{total}_{\mathrm{speed}}}$$

In the case of Virgo, the vacuum enclosure has a wall surface above 25000 m^2. The only way to keep the pumping system at an economically affordable level is to reduce the wall outgassing rate. While designing Virgo, available experience on very large systems was that the best outgassing rate reachable, after bake-out, was of the order of 10^{-10} Pa $\cdot$ l / cm^2 $\cdot$ s. Before Virgo and LIGO construction, the largest existing UHV systems were at least one order of magnitude smaller. The outgassed gas species are largely dominated by Hydrogen, trapped inside the whole sheet thickness in the stainless steel production process, at a concentration of a few ppm in weight. With the given values the total pumping speed should be 10^6 l/s, with a moderate safety factor of 4. The cost of a similar pumping system, distributed along a 6 km length would have doubled the Virgo construction budget.

2012 JINST 7 P03012

The problem was solved by experimenting with and developing the air-firing treatment on material samples and on several tube and tower prototypes. Firing involves heating the raw material or the already built tank elements at several hundreds of degrees Celsius, in order to increase Hydrogen atom mobility and rapidly reach the equilibrium with the much lower H_2 concentration in the atmosphere. After cooling down, Hydrogen concentration in the metal sheets and outgassing rates are reduced by more than two orders of magnitude. Air-firing is a cheaper alternative to traditional vacuum oven-firing, which is not applicable to very large systems.

For Virgo, a very large oven was built by the CNIM company (La Seyne sur Mer, Toulon), capable of reaching 450 °C and meeting the requirement of an H_2 outgassing rate well below $5x10^{-12}$ Pa $\cdot$ l / cm^2 $\cdot$ s. Similar results were obtained for the large valve components and for the tower lower vacuum tank, fired in the same oven for a longer time, given the larger thickness of the metal sheets to be Hydrogen depleted. The firing temperature was chosen to be well below the brittle temperature for 304L stainless steel.

The systematic study on stainless steel Hydrogen outgassing, in addition to the development of the air-firing technology, allowed us to obtain relevant vacuum technology results [13–15]. In particular, we demonstrated that, for our commercial cold-rolled austenitic stainless steel sheets (Avesta, Sweden), diluted Hydrogen is really desorbed by air-firing while surface oxide layers do not play a significant role in reducing outgassing [16].

Air-firing and 150 °C bake-out, together with a strict application of standard washing procedures also allowed us to meet the hydrocarbon maximum pressure requirement.

4.2 The Towers

4.2.1 Function and specifications

The vacuum towers have been conceived to host, in vacuum, the interferometer optical elements and their anti-seismic suspensions, the Superattenuators. The various optical elements require seismic isolation at different levels, hence the SA have different heights. The main mirrors (Beam Splitter, Power Recycling, North Input, North End, West Input and West End), acting as gravitational test masses, have long suspensions inside 11 m-high towers (figure 16).

The Mode Cleaner mirror, the Input Bench and the Detection Bench have short suspensions and 6 m-tall towers (figure 17).

The mirrors contained in the long towers require an ultra-high vacuum level (10^{-7} Pa), not compatible with the electromechanical elements of the suspensions, from which they are suspended by means of metallic wires. Hence the long towers are vertically split into three compartments: the lowest section (10^{-7} Pa) containing the mirror 'payload' and the propagating light beams; the upper part (10^{-4} Pa) containing the anti-seismic suspension; and the intermediate compartment, needed to obtain, by differential pumping, the appropriate pressure difference between upper and lower compartments. The three volumes are pumped independently and are only connected by a 16 mm hole, required to enable passage of the payload suspension wire. In order to keep the conductance between the three compartments as low as possible, the holes are equipped with 200 mm long 'conductance tubes'.

UHV conditions in the tower lower part can be reached by respecting strict cleaning conditions, followed by a 150°C bake-out process. To this end the long tower bases have been enclosed inside appropriate ovens.

2012 JINST 7 P03012

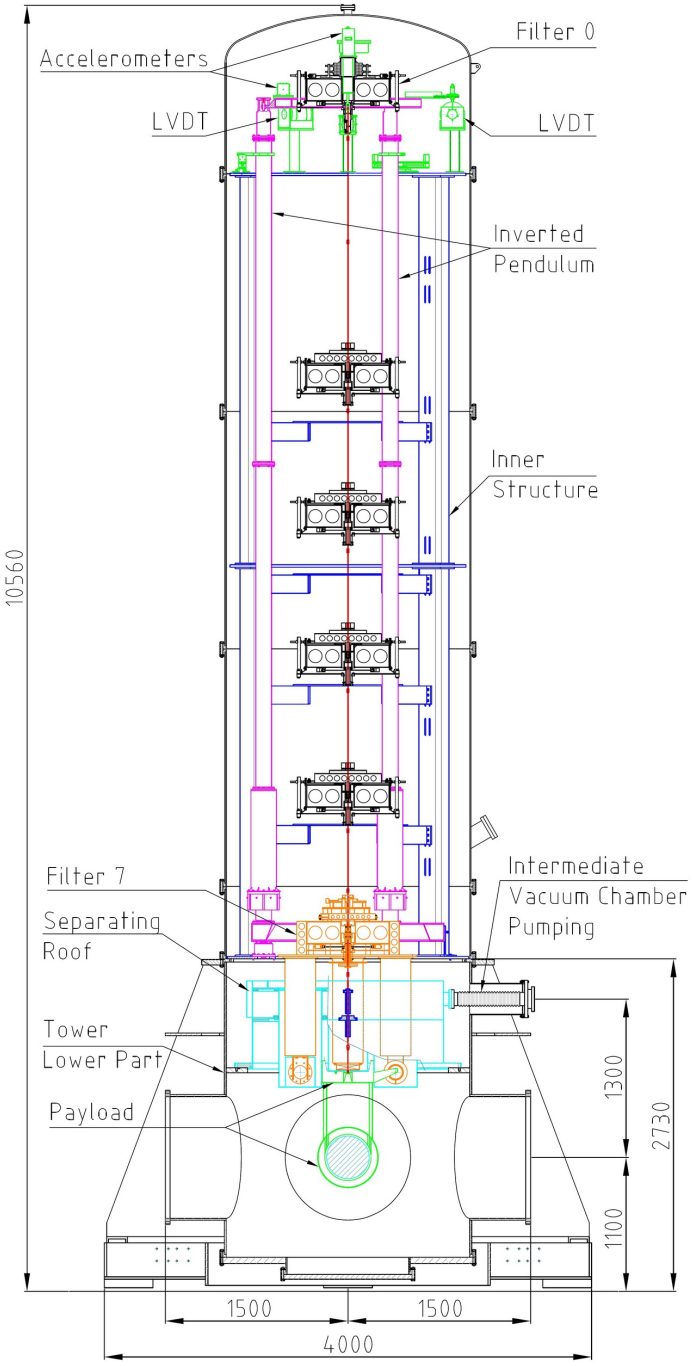

Figure 16. Vertical cross-section of a long tower, showing inner details and the SA.

The MC mirror and the optical benches require only a high vacuum level (10^{-4} Pa), hence the short towers have one single compartment and no ovens.

Only the base of a tenth tower has been installed between the Beam Splitter and the Detection Bench. This is provisional for the installation of a Signal Recycling mirror and its corresponding suspension.

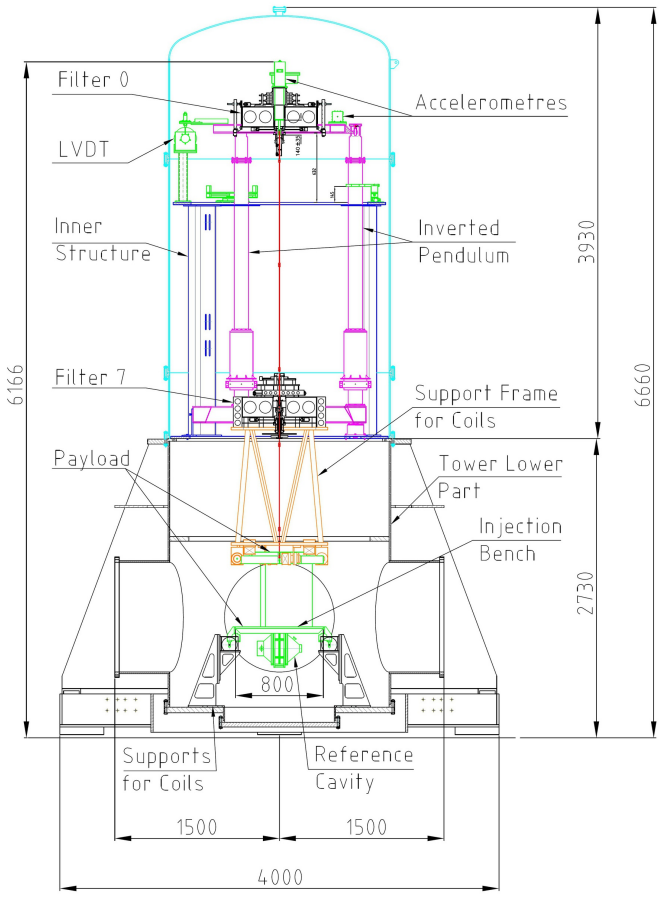

Figure 17. Vertical cross-section of a short tower, showing inner details.

The tower bases are linked by vacuum tubes for the interferometer beam propagation; the link tubes between the central towers are 300 to 400 mm in diameter; the arm tubes, containing the 3 km-long Fabry-Perot cavities have a 1.2 m diameter (figure 15). The link tubes are equipped with suitable aperture gate valves, in order to allow independent evacuation of the towers. Link tubes connecting towers with different working pressure are closed by vacuum-tight optical windows, set at the Brewster angle with respect to the light beam, in order to reduce unwanted reflections.

Tower bases are equipped with several optical window ports, for the passage of auxiliary laser beams, used for mirror position feedback control and for monitoring purposes.

4.2.2 Mechanical design

A unique feature of Virgo is its operational sensitivity, starting at a frequency as low as 10 Hz. This is achieved with the SA, which is already able to reduce the seismic motion of the test masses (the mirrors) at that frequency. As a consequence, all the resonant frequencies of the mechanical structure supporting the SA have been pushed to above 10 Hz, i.e. well above the SA cut-off frequency (a few Hz). A tower design in agreement with these conditions was achieved by detailed modelling via a simulation code. In addition, the SA structure, described in a following dedicated section, was

2012 JINST 7 P03012

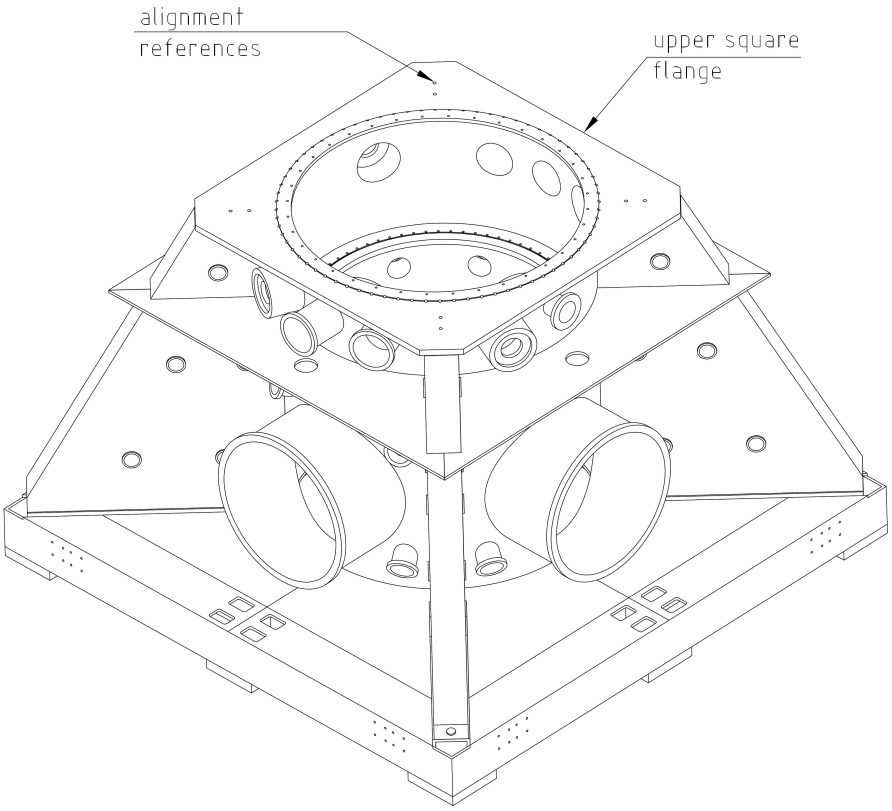

Figure 18. A tower base.

taken into account. The mechanical simulation was performed in all possible conditions, including bake-out under vacuum. The lowest collapse factor turned out to be 12.8.

The mechanical structure of the tower consists of a stiff base, allowing for an accurate positioning and a strong clamping to the building floor, and of an upper part, made of modular cylindrical pieces. All pumping ports, feed-through and windows are located on the base; so no equipment has to be disconnected when opening the tower.

4.2.3 The tower base

The tower base hosts the optical payload and serves as a high stability support for the SA. The base (figure 18) consists of a cylindrical tank (2000 mm diameter, 2740 mm height, 15 mm wall thickness), welded to an upper square flange (2400 mm x 2400 mm x 60 mm), which is taken as a reference for the tank geometry and for the tower positioning. To this end, four precision holes have been drilled at the flange corners. The tank is welded to a very massive (4000 mm x 4000 mm x 300 mm) square pedestal clamped to ground. Pedestal, tank and upper square flange are bound with four vertical triangular stiffening wings.

Each tank has four large perpendicular ports, in the horizontal plane, at the interferometer beam level, i.e. 1100 mm above the building floor. All the ports have a 1000 mm diameter, except the ports connected to the arm tubes, having a 1200 mm diameter. The 1000 mm ports are closed by flanges equipped with 400 mm ports for the passage of the main interferometer beams, through

2012 JINST 7 P03012

the link tubes. The pumping ports for the three tower compartments are all located on the upper part of the tower bases, just below the square flange. Several smaller ports and optical windows are present for the main laser beam and monitoring and control beams.

The bottom walls of the towers consist of a circular 2 m diameter lid, bolted under the tower to an appropriate flange. In the centre of the lid is a 1 m diameter aperture, which is closed, in turn, by a secondary lid. The 1 m aperture allows the installation of the payload from below, through the clean gallery existing under the floor of the buildings. In exceptional circumstances, the 2 m lid might also be opened.

The bases of all the towers are structurally identical; any differences only concern optical windows and access port positions. All of the large ports (1 m, 1.2 m, 2 m) are closed by bolted flanges and metal gaskets (Helicoflex Delta, by Cefilac, France); this is particularly important for bakeable towers. Only the lower 1 m lid has a temporary Viton gasket, given that it is opened relatively often in order to enable work to be undertaken on the payloads. The 400 mm ports connecting the tower bases through the link tubes, have bolted flanges with Helicoflex gaskets. Ports with a diameter of 250 mm or smaller have standard Conflat flanges and copper gaskets.

4.2.4 The Separating Roof

The two horizontal disks vertically dividing the long towers into three volumes with different vacuum levels, constitute one single rigid block called the 'Separating Roof'. This system is located inside the tower base, just below the upper square flange and rests on an appropriate inner flange. Besides the limited gas flow through the low-conductance tubes (2 l/s each, combined in series), the vacuum tightness between upper and lower compartments is reached without gasket thanks to the very high quality machining of this flange and of the lower surface of the Separating Roof. The two surfaces can slide, one over the other, without loosing tightness. This feature allows the adjustment, within a 25 mm radius, of the Separating Roof position on the horizontal plane, following the adjustment of the mirror suspension wire, which is operated upon at the inverted pendulum top stage. Minor adjustments, up to a few mm, can be accommodated inside the conductance pipes; larger adjustments require the movement of the Separating Roof. This operation is achieved from outside, by a vacuum-tight system of micrometric screws and bellows. The wire position inside the conductance pipes can be visually inspected by a set of mirrors and a TV camera (periscope).

The Separating Roof structure is complicated by a number of features. The last SA stage, suspended above the roof has four legs supporting the coils acting on the magnets of the 'marionette', the underlying SA stage (see Chapter 6), located under the roof. To keep the upper and lower tower compartments separate, the Separating Roof is traversed by four cylindrical pots enveloping the coil legs. At the end of each pot is a thin Kapton window, enabling steering coil/magnet interaction, without magnetisation effects. Pumping and venting speeds have been dimensioned in order not to break the thin windows as a result of too-high pressure differences (below 600 Pa).

The marionette supporting wire, passing through the conductance tubes and coils and magnets located on opposite sides of the Kapton windows, demands an extremely difficult sub-millimetric positioning of all of the parts. All of this does not hold for the short tower, however, which is not equipped with a Separating Roof.

4.2.5 The upper cylinder

Above the base there is a 2 m diameter, 40 cm high cylinder, known as the 'technical ring', with four 250 mm ports, each equipped with all of the required electrical feed-throughs for in- and outgoing signals.

Above the technical ring, the tower tank is completed by three (one for the short towers) standard cylinders and a top cap. The 2 m flanges coupling the cap, the standard cylinders and the technical ring to the base, have Viton gaskets, not being part of the UHV compartment, and are hence not included inside the bake-out oven.

4.2.6 Construction

The towers were built by the company SDMS (St. Romans, Grenoble). Fabrication took place in a clean environment, free of dust and of oil vapours, in order to meet outgassing and cleanliness requirements. To the same aim, the inner surfaces of the tanks have been electro-polished (RA 1.6 μm, rugosity, following international mechanical standards). The different tower parts were processed by thermal treatments in air, up to 450 ^{o}C, and under vacuum, up to 200 ^{o}C, for mechanical tension relieving, Hydrogen outgassing and accurate de-greasing.

At the end of the fabrication process, each tower base was carefully surveyed, using, as reference, the four precision holes on the upper square flange. Once cleaned, the base tanks were sealed and conditioned under vacuum until their installation on site.

4.2.7 Cleanliness

As explained in the requirements section, it is extremely important to protect the optical payloads against both hydrocarbons and dust. Having done this during tower fabrication, cleanliness may be endangered when the towers are opened for SA or payload installation.

To prevent this, air conditioning and clean air circulation circuits are arranged in order to have slightly increasing pressure in the following series of volumes: the external atmosphere, inside the Central and Terminal Buildings, the clean rooms and the gallery, inside the tower base. The environment becomes cleaner and cleaner moving in the same sequence, up to class 10 inside the towers, in the absence of operators.

To prevent dirt entering the towers, bases are equipped with two 200 mm pipes injecting, just below the Separating Roof, 1200 m^3/h of filtered air, derived from the main clean room air system. In this way, when the tower base is opened from above, for SA installation, an over-pressure is formed in the UHV payload compartment and clean air flows upwards, through the conductance pipe, preventing the entrance of dirt. The Separating Roof acts as a sort of umbrella against dust. Moreover, the work volume above the tower is confined by a plastic tent, enveloping the whole service scaffolding.

When the tower is opened from below, for payload installation, the filtered air flow acts as a real air shower, keeping clean both components and operators entering from the gallery below. In order to establish a controlled circulation of clean air, appropriate apertures, close to the gallery floor, inhale air returning it to the filter units. The established regime is not at all a laminar flow, due to the insufficient inlet aperture and to the relatively large dimensions of payload and operators, with respect to the UHV chamber. As described in the payload section, dirt accumulation is monitored with silicon wafers located in the most exposed points of the chamber.

2012 JINST 7 P03012

4.2.8 Link tubes

The link tubes are the 400 mm vacuum tubes connecting the towers in the Central Building to allow the passage of light beams. The link tubes have to account for:

- their own dilatation and tower dilatation due to 150°C bake-out;
- misalignment between tower bases;
- mechanical tolerances on tower flange and link tube construction;
- installation of tower/tower sectioning gate valves;
- installation of stray light-absorbing glass baffles.

Hence a link tube, up to 3 m long, includes an elastic bellows and a 250 mm aperture gate valve. The bellows allows the elastic longitudinal displacement of $\pm$40 mm and transverse displacements of $\pm$15 mm. The whole link tube, including valve and glass baffles, is bakeable. All the gaskets used for the link tubes are metallic, apart the valve gate sealings, which are Viton, and bakeable up to 200°C. The link tubes also serve as connections to the permanent pumps for the tower UHV chamber. Such permanent pumps are constituted by two large tanks with Ti evaporators, coupled with an Ion pump; given their large dimensions, these groups cannot be installed directly on the towers.

4.2.9 Bake-out ovens

In order to be able to bake-out the tower bases, they have been enclosed inside large hot air flow ovens, integrated with the lower part of the scaffolding surrounding the towers. The upper part of the tower bases, where all of the pumping ports are present, is outside the oven; it is thermally insulated with soft blankets, in order to reach a sufficient temperature on all the walls of the UHV chamber. The structure is permanent, while the insulating walls and the heating fans are movable. In each tower the available power is 20 kW, distributed by a circulating air flow of 2000 m^3/h. Under the tower there is a 3 kW heating disk for the tower bottom lid. The power supply and control unit installed in the Central Building are able to power two tower ovens at a time.

The ends of the link tubes, including the gate valves, are inside the oven volume; the tube external part is heated by tape resistors and covered by insulating blankets.

4.3 The arm tubes

4.3.1 Vacuum performances and related constraints

As stated in the requirements section, the maximum tolerable pressure in the arms is 10^{-7} Pa. This low pressure requires a very small outgassing rate of the tube material, which should be $< 5{\cdot}10^{-12}$ Pa·l/s·cm^2 for Hydrogen, and $< 5{\cdot}10^{-16}$ Pa·l/s·cm^2 for hydrocarbons; this cleanliness, together with the absence of Fluorine, being mandatory for the avoidance of pollution and chemical attack of the mirror coating.

The design final pressure (10^{-7} Pa) is reached by a bake-out of the whole evacuated arm tube at 160°C, in order to remove water vapour.

2012 JINST 7 P03012

Finally, the gas flow due to leaks has to be well below the gas flow due to wall outgassing. Hence the maximum acceptable leak rate for an assembled arm tube has been established at 10^{-5} Pa·l·s^{-1}, summing up the contribution due to the tube modules, as built in the factory, and the contribution due to the circular welds performed on-site to assemble the modules. With 200 modules per arm, each module cannot leak more than $3{\cdot}10^{-8}$ Pa·l·s^{-1}.

4.3.2 Mechanical design of the arm tubes

The main mechanical constraints retained for the tube design are:

- resistance to a 3 bar external pressure (operate in vacuum with a safety factor of 3);
- remain within the elastic domain for a temperature between 5 and 160°C (bake-out);
- axis of each module of the 3 km tube inside a 10 cm diameter cylinder;
- accept a vertical shift up to 1 cm between adjacent supports;
- have the possibility to remove and replace an already installed tube module;
- the presence of baffles inside the tube to prevent scattered light.

Two completely different concepts have been considered to cope with these requirements: plain tube pieces inter-leaved with bellows and a completely corrugated tube.

In addition to the choice of a mechanical principle for the tube, there were also geometrical constraints, one of which related to the choice of having the possibility to host 3 interferometers, which implied a minimum diameter of 1.2 m, the other being a realistic length of tube prefabricated modules, in order to minimise the work required to weld the modules together, while being transportable by truck.

4.3.2.1 The basic module

After the construction and test of several prototypes, both straight tube and corrugated modules, it was decided to build 15 m long stainless steel elementary modules with a diameter of 1.2 m, outer stiffening rings and one bellows per module, allowing for dilatation. Stainless steel ‘lips’ at both ends would allow weld together of the modules and easily separate them, cutting out the welded portion of the lips.

The different components, and then the complete tube, were modelled and optimised with SYSTUS, a finite elements programme, [17] determining the following parameters for the basic components:

- wall thickness for the tube: 4 mm;
- distance between 2 stiffeners: 1.2 m;
- bellows wall thickness: 2 mm, 10 waves with a shape such as that shown in figure 19 (Bezier shape);
- lip thickness: 2 mm, with a shape such as that shown in figure 19;

2012 JINST 7 P03012

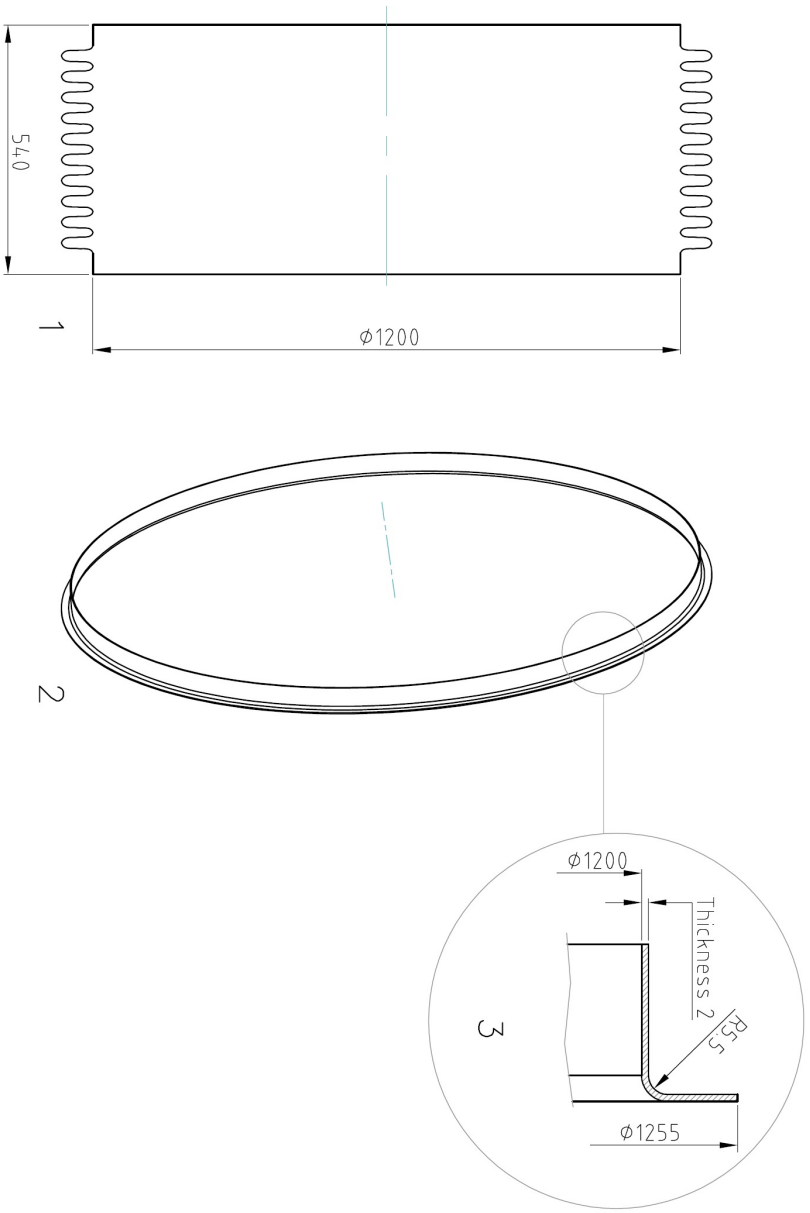

Figure 19. Bellows and lip shape.

- defect of circularity of the external edge less than 0.5 mm for assembling two modules with a welding robot.

Once welded together, it must be possible to disassemble the modules by nibbling the welded lips with a cutting robot. Since it was decided that 3 nibblings should be possible, the height of the lips was fixed at 22.5 mm.

4.3.2.2 Composition of one arm

A complete 3 km tube is composed of 300 m-long basic sections, each with 20 modules, the first and last of which are respectively a 'reinforced module' and a 'pumping module' connected to a pumping station. Each arm tube is composed of eight such sections and, at both ends, by two 'end sections', which are also 300 m-long and which include large valves and special short modules to link the valves with the towers (figure 20).

Each section was planned to be vacuum tested independently, and was closed provisionally by welded curved end caps, when being tested. So both end surfaces withstand the atmospheric pressure (12 t force) and transmit it longitudinally via the module feet to the tube supports, which are rigidly fixed to the tunnel floor. This implies that the two end modules and their supports, of each section, are especially reinforced.

This leads to several types of modules:

- basic module: (figure 21), 15 m-long, made of 3 stainless steel (4 mm thick) plain cylinders welded together, with stiffening rings spaced 1.2 m apart. At one end of the module is a 54 cm-long bellows, with 10 68 mm-high waves, and a lip; on the other end the module is terminated by a simple lip;

- reinforced module: essentially the same as the basic module, but without bellows and with more stiffening rings (one every metre).The absence of the bellows leaves a free space between two sections to install, then remove, the end caps for the vacuum test of the sections. The bellows were installed later, in order to join the tested sections and compose a 3 km arm tube;

- pumping module: the same as the basic module, but with 3 pumping ports: two DN200 for intermediate and permanent pumping, one DN63 for measurement instrumentation;

- link modules: (ϕ 1200 mm) of different lengths to connect tubes, valves and towers. When necessary the link modules contain a bellows.

From the previous description, it turns out that the large majority of the modules building a 3 km arm tube are connected by welds. This is, by far, the most reliable and economically convenient solution. Only the connections of tubes to towers and large valves are made by bolted flanges, coupled by metal gaskets. For a maximum tightness and reliability, Helicoflex gaskets have been used. On the large valve side connected to the arm tube, the metal gasket has also been doubled by a Viton gasket, to facilitate leak search.

The modules rest on supports fixed to the tunnel ground, by means of stainless steel feet (figure 21) welded on to the last stiffening rings on each side of the basic module, or to special double rings at one end of the pumping module and at one end of the special reinforced module. On these 2 last modules (reinforced and pumping), the feet transfer the 12 tonne longitudinal force, due to atmospheric pressure, from the tube end cap to ground through the support, which has to be especially strong.

2012 JINST 7 P03012

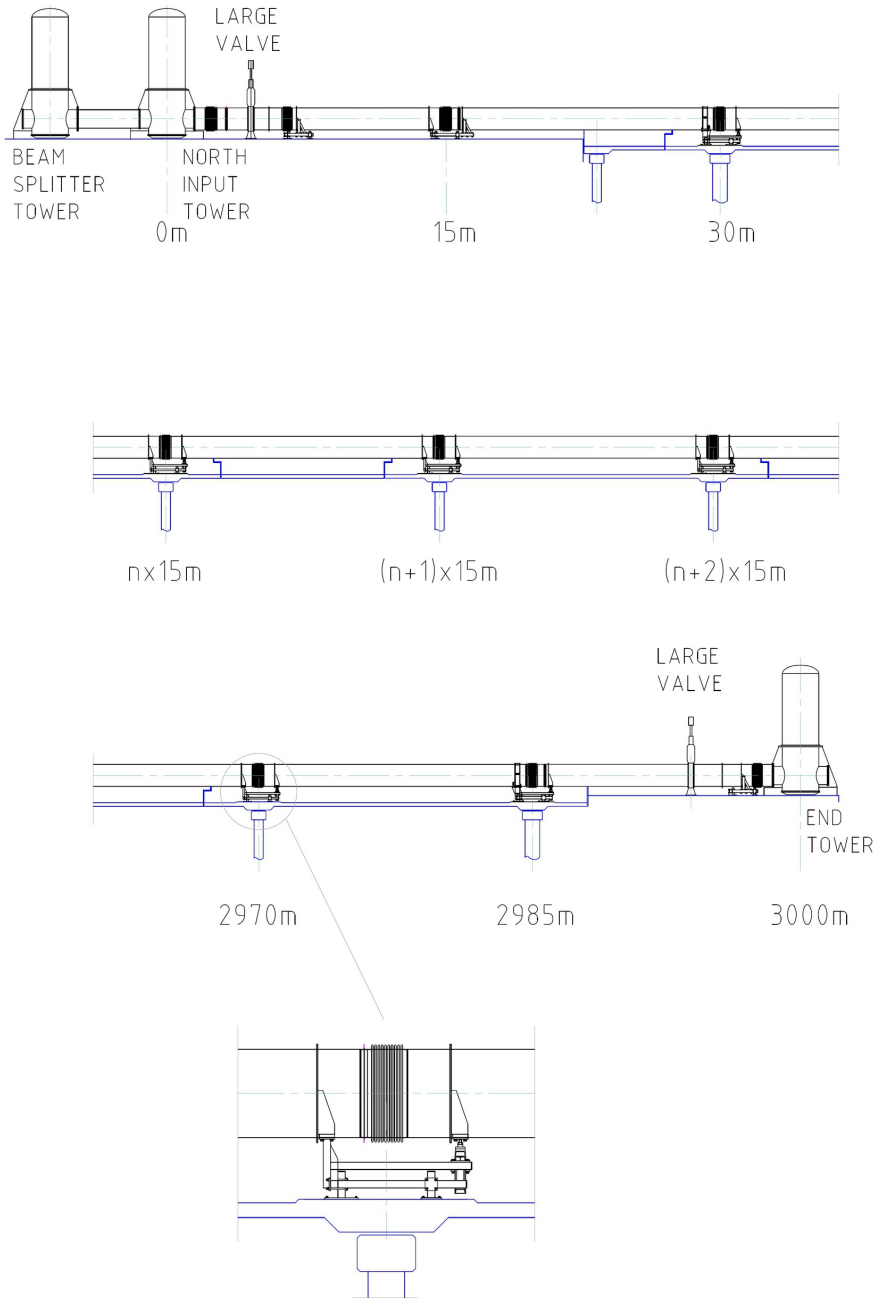

Figure 20. Composition of an arm tube, showing the first modules, some intermediate modules and the end modules.

4.3.2.3 Module construction

For the construction of the 404 modules, a call for tender was launched in 1997, and the contract signed with Société CNIM from La Seyne-sur-mer (France); the raw stainless steel (1.4307, equivalent to 304L) was supplied by Avesta (Sweden). CNIM developed a specific industrial tool, with huge benches for: the cutting of the stainless steel plates; rolling them into cylinders; longitudinal

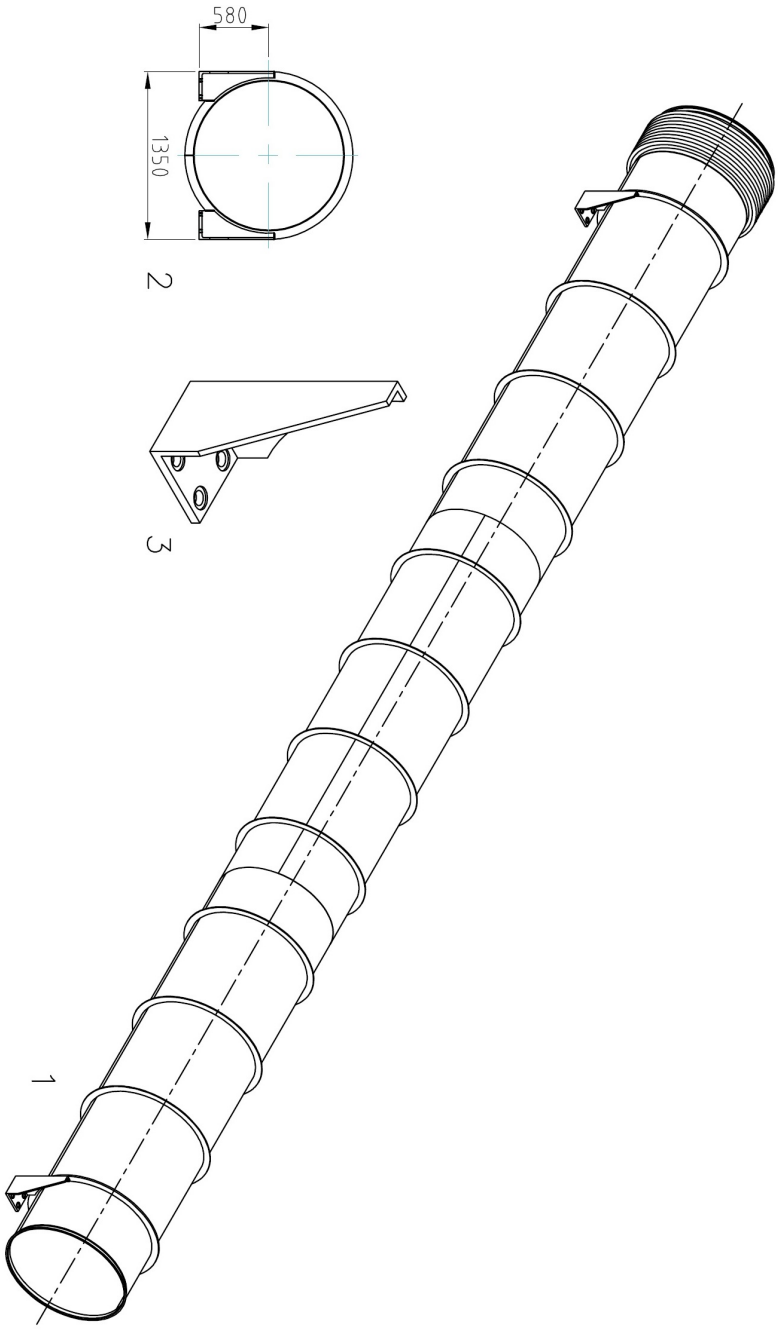

Figure 21. Basic module, its cross-section and one foot.

welding; circular welding for assembling the different parts of a module (figure 22), and a specific tool for precisely installing the stainless steel baffles (see the dedicated chapter). All this was done while avoiding contact between the module inner surface and materials other than clean stainless steel. Subsequently, in a 'white' hall (class 10000), the modules were washed with alkaline soap in warm water, then rinsed with de-mineralised water, in a 17 m-long washing machine. After drying

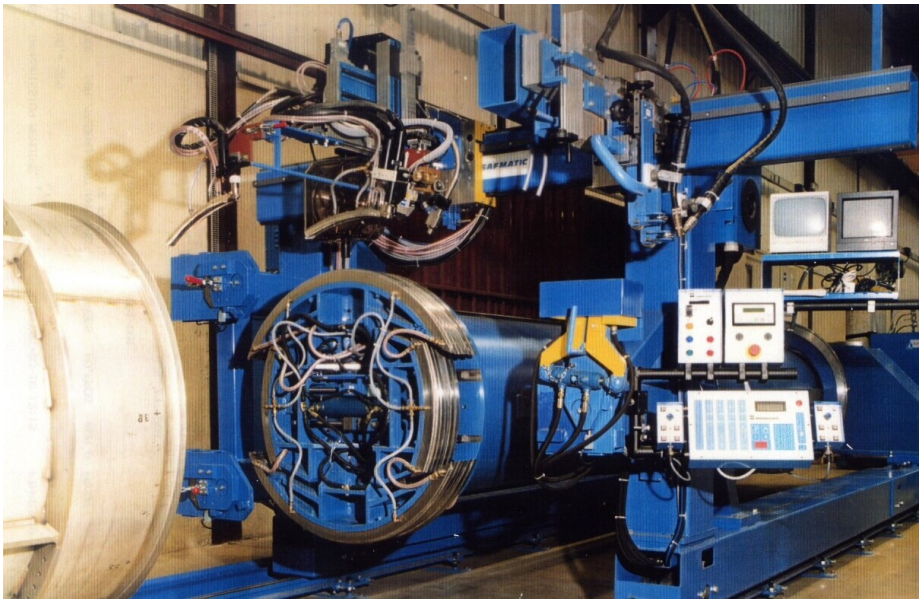

Figure 22. Picture of CNIM circular welding machine.

they were geometrically checked on a built-for-purpose reference bench. A big oven (17·5·5 m^3) with electrical heaters and two permanent independent airflows (to remove the desorbed Hydrogen inside and outside the modules), made it possible to heat 4 modules together at 400°C for 96 hours, necessary for Hydrogen desorption and residual hydrocarbon evaporation. Finally, on a Helium test bench, the tightness of the module was tested, together with its cleanliness, by Residual Gas Analysis (RGA).

The bellows, as stated above, were optimised in the complete model [17], for a maximum compression of 37.8 mm between 5°C and 160°C, while remaining in the elastic domain.

These were 'hydro-formed' at the SFZ (France) firm and tested, then heated to 400°C at CNIM, before the stiffness of each of them was measured to ensure it was within the expected range. After the 400°C heating, the length had increased by a small value (up to 3 mm) but was a little different for each bellows. In order to keep the module length tolerance (15000 $\pm$1 mm), CNIM had to build the straight part of the tube adapted to the actual length of the associated bellows. After the final 400°C heating no change in length was detected.

Lips were also built at SFZ, with the right tolerances, which were not retained following the TIG (Tungsten Inert Gas) welding on the modules. Therefore, in order to keep the required tolerance (0.5 mm on circularity), these had to be re-machined once welded on to the modules.

The whole industrial process was validated with a first module ('module 0'), which was fully tested, including bake-out at 160°C to eliminate water vapour. The results showed that the obtainable pressure limit was correct. The fact that hydrocarbon rates were slightly too high signified that new efforts for improving the cleanliness throughout the whole process were required, then production could begin.

Eventually, the mass production of the modules (404 modules or links) lasted from 1999 to 2002, with a maximum production rate of 20 per month. For each module the results of the tests

2012 JINST 7 P03012

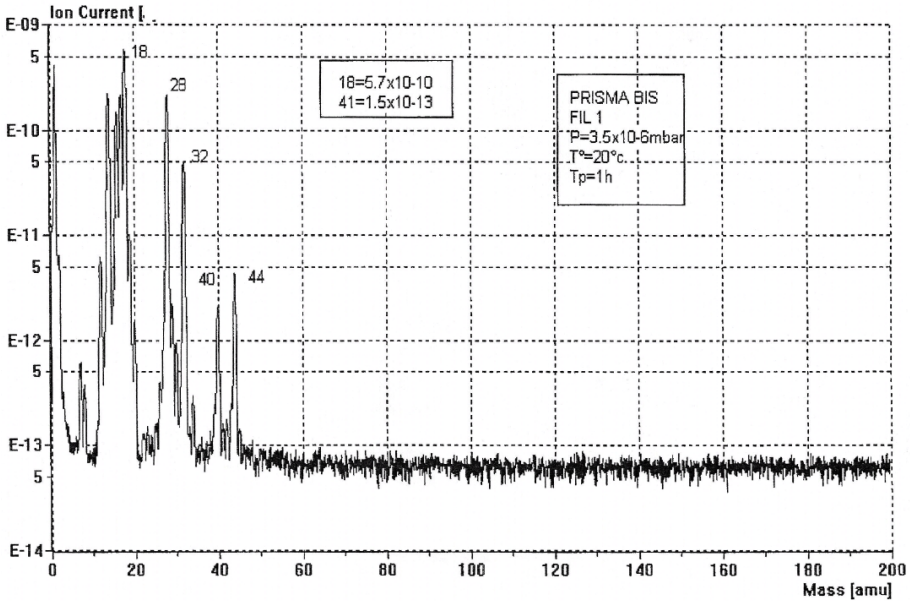

Figure 23. Typical residual gas spectrum for a finished tube module.

are part of a report, summarising all of the geometrical measurements, tightness results, cleanliness (figure 23 shows the result of an RGA analysis), and data on the stainless steel quality.

Following the final tests, the modules were closed with tight stainless steel end disks; the inner volume remained connected to the external atmosphere by a small hose through a Teflon/active carbon filter, in order to keep the inside of the module in equipressure with the outside (breathing modules). Then the modules were carefully packed to be transported by truck or train.

4.3.3 Thermal insulation

Before going to Cascina, the modules were shipped to the Belleli premises in Mantova (the company in charge of the on-site assembly) in order to be equipped with a thermal insulating cover. This insulation is necessary in order to minimise the heat-loss when heating the tube with an electrical current during the bake-out process. The cover installed on the rigid part of the module is made of two over-lapping rockwool layers (with a total thickness of 160 mm) covered with an aluminium sheet to provide a high uniformity temperature distribution. After assembly on site, the bellows were also covered with insulating mattresses, in order to complete the insulation. With this equipment the average power required to obtain a 150°C temperature is about 200 W per metre of tube.

At the same time thermocouples were installed on the module body in order to control the temperature during bake-out.

Following this installation, the modules were shipped to Cascina by truck, where they were assembled to compose the arm tubes.

2012 JINST 7 P03012

4.3.4 Bake-out system

The tubes are heated by Joule effect, sending a DC current through the electrical resistor constituted by the tube wall itself; this is possible since every arm tube is completely welded over 3 km and the tube feet are electrically insulated from supports and ground by fibreglass washers. The DC current choice was necessary for safety reasons and to avoid unwanted electromagnetic inductance effects (high impedance, uneven current density distribution). Current is supplied to the tube at five points (300 m, 900 m, 1500 m, 2100 m and 2700 m) by AC/DC converters. Intermediate points (0 m, 600 m, 1200 m, 1800 m, 2400 m and 3000 m) are grounded. In this way the maximum voltage is safely limited to 55 V and the current (about 1700 A) flows from each supply point, in opposite directions, to the two next ground points. Power is supplied by diesel generators, to save the high cost of two permanent supply lines, 3 km-long, each with a total power above 0.6 MW. The diesel generators are rented, as the bake-out operation only lasts a few weeks per arm and is very seldom performed (only once, up to now). The AC/DC converters have been purchased, being custom-designed and built; as one arm tube is baked at a time, six converters have been bought (five plus a spare). Every 300 m a suitable copper current connection is bolted to a tube stiffening ring; these are alternatively connected to ground and to 55 V. The return current conductor is constituted by a 10x200 mm^2 pure Al bar, attached to the tunnel wall and having a resistance five times-smaller than that of the tube.

4.3.5 Tube supports

The functions of each support [18] are:

- bear the ends of two adjacent modules allowing only longitudinal relative displacements (Z direction) due to thermal expansion, exploiting the elasticity of the intermediate bellows; relative transversal displacements (X, Y) of the module ends have to be strictly prevented to preserve the lip weld integrity;
- hold the forces coming from temperature variations of modules and from vacuum;
- allow for the alignment of the modules in the X-Y plane (perpendicular to the tube axis) referred to the given reference line;
- allow for the fine adjustment of the module position in the Z direction referred to the reference mark on the ground of the tunnel;
- during assembly, allow the fine adjustment of the modules to fit two adjacent lips in the same circumference; to allow correct lip welding;
- allow the re-alignment of the modules after installation if it is necessary to compensate for ground subsidence.

The carrying capacity is determined by the fact that a standard module has a maximum weight of about 3400 kg, including thermal insulation, and a pumping module has a maximum weight of about 3800 kg, including thermal insulation and vacuum equipment.

One further requirement for the supports, is that they should not induce any vibrations on the tube when accommodating movements due to temperature changes, pumping etc. So rolling or

2012 JINST 7 P03012

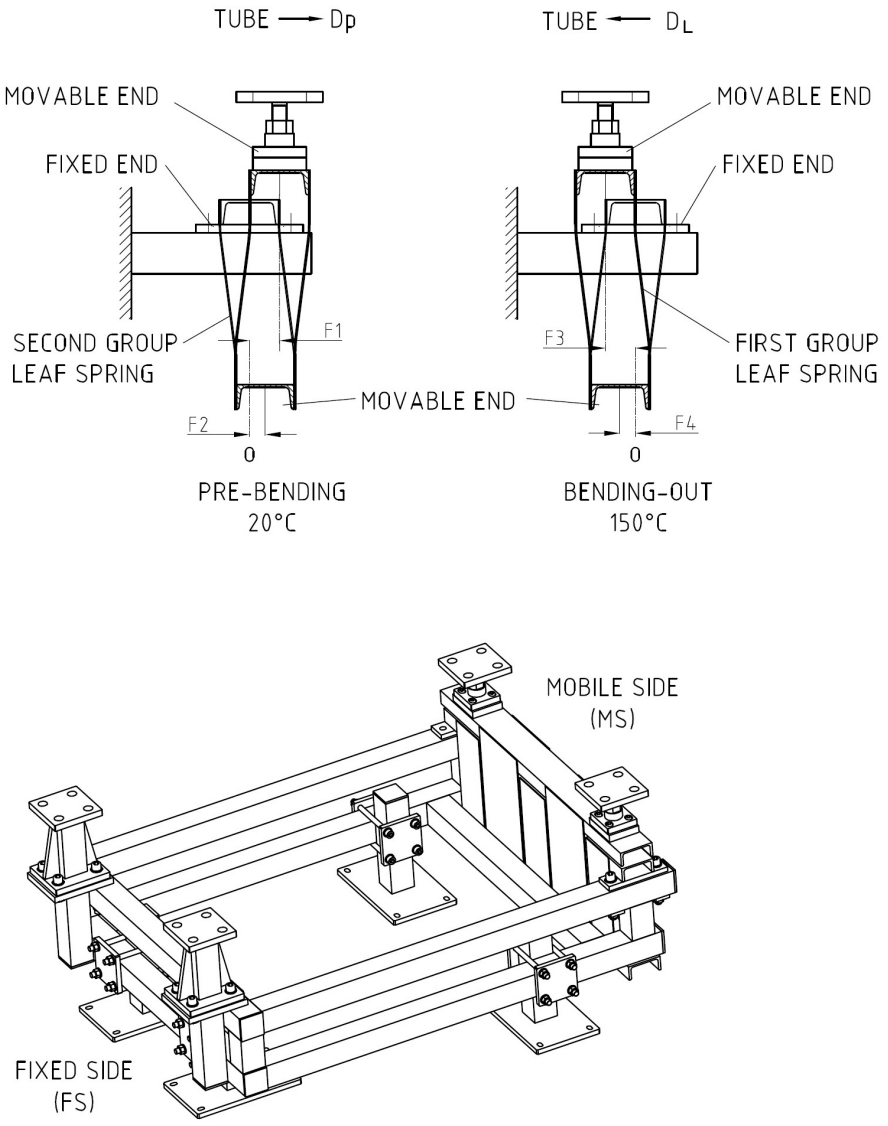

Figure 24. A basic support and the principle of its mobile side.

sliding systems were rejected, and the chosen solution was leaf springs, which allow displacement by bending 'elastic' steel plates.

4.3.5.1 Basic support

A view of a basic support is shown in figure 24.

The basic support consists of a strong steel frame resting on four adjustable legs, which are bolted to the tunnel floor. On top of the frame are four plates, supporting the two feet at the end of one module and the two feet at the end of the adjacent module; in the region between the two feet pairs are the expansion bellows and the lip weld joining the two modules.

Each support has a fixed side (FS, rigidly connecting the module feet to ground) and a mobile side (MS, connecting the module feet to ground, but allowing longitudinal displacements only).

The fixed parts are blocked to ground and the 2 feet of each module on the side without bellows are bolted to this fixed part. The other part, which receives the feet on the bellows side of the adjacent module, is mobile, in order to allow for the 37 mm dilatation of the module (from 5 to 160°C). At module installation, the position of the FS is fixed first, while the position of the MS has the possibility to match the module feet with their fabrication tolerances.

Before performing the lip welding, a fine adjustment to fit the two lips is performed.

Figure 24 also shows the cross-section of the mobile side of a support, based on steel leaf springs. In order to allow only parallel (horizontal) displacements of the plates supporting the module feet, leaf springs have always been used in pairs, welded to an upper and a lower U-beam. The use of two sets of leaf springs in series, one set going down and one set coming up, reduces, virtually to zero, vertical displacements of the upper mobile U-beam; unimportant vertical displacements can only occur in the lower U-beam. At module installation the leaf springs were bent in order to match them to the length of the module at the actual temperature. The use of tools and clamps with precision screws allowed for the adjustment of module ends with sub-millimetric accuracy, before lip weld execution.

The adjustment ranges allowed by the supports in X, Y and Z directions are shown in figure 25.

4.3.5.2 Special supports

Every second basic support was slightly modified in order to obtain some mobility (a few mm) also on the 'fixed side' (FS), introducing short leaf springs. This was necessary in order to compensate for the thermal dilatation of the tunnel concrete beams. These supports are called 'balance supports'.

There are also 'Superfixed' supports for the ends of the 300 m sections, which have to stand an axial force of ~ 12 tonnes, because of the atmospheric pressure, during the preliminary vacuum test. These supports are reinforced with strong, inclined beams.

Other special supports [18] had to be built for link sections near towers and valves.

4.3.5.3 Electrical insulation

Since the bake-out of the tube is done by DC current flowing through the tube (U_{max} = 60 V), the tube must be electrically insulated from the ground. Insulation fibreglass blocks have been inserted between support upper plates (made of carbon steel) and module feet.

4.3.5.4 Construction

The supports were built at the firm Belleli, who also ensured their positioning and mounting in the tunnel.

4.4 Mode Cleaner tube

The Mode Cleaner tube links the Injection tower with the Mode Cleaner tower. These three items constitute a vacuum chamber separated by a window from the rest of Virgo, the required pressure (10^{-4} Pa) being not as low as in the Virgo tubes and central towers. The Mode Cleaner is a resonating optical cavity conceived to reduce sources of laser noise, such as acoustic waves, thermal gradients on the optical path, motion of dust, etc.

2012 JINST 7 P03012

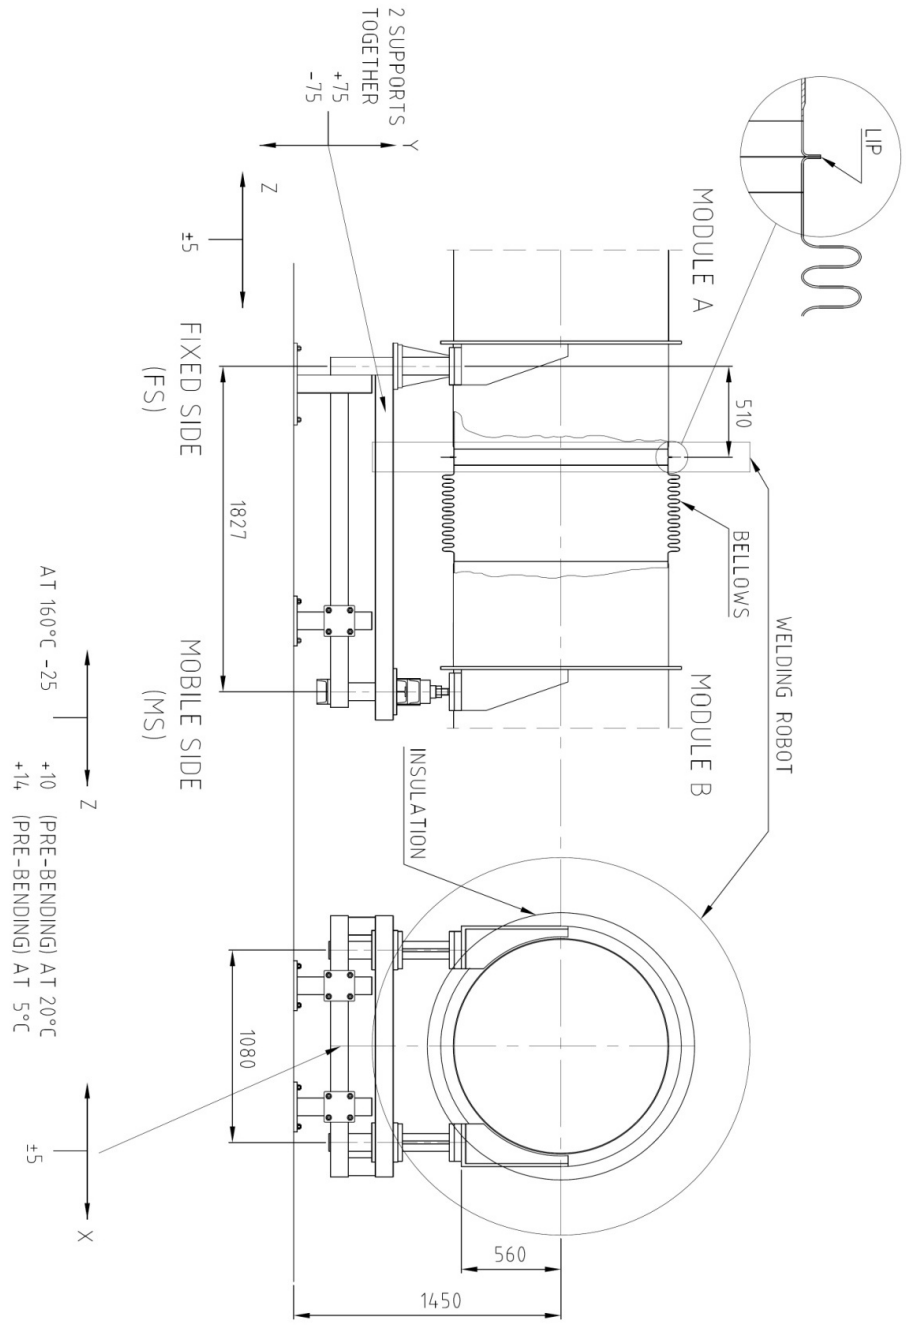

Figure 25. Support adjustment ranges in X, Y and Z, expressed in mm.

The tube is 142 m-long and 300 mm in diameter, with a stainless steel (1.4307) wall thickness of 3 mm. It is composed of 25 modules, 5.5 m-long. Four bellows are spread out over the whole length, to obtain a working temperature between -10° C and + 40° C. Except for 2 modules on each end of the tube, modules and bellows are TIG welded together. Close to the Injection Tower a DN 250 valve allows for the isolation of the Mode Cleaner; this is followed by 2 modules, mounted

with bolts and Viton gaskets, having two ports, one CF63 for measurements and one CF100 for an eventual additional pumping group, the tube being normally pumped via the Mode Cleaner Tower. The two modules close to this tower are also mounted with bolts and viton gaskets; one of them enables the adjustment of the tube length.

The tube rests on 25 adjustable supports (3 fixed and 22 sliding supports). The three fixed supports, together with the two towers, define four independent sections with one bellows each. The tube has been aligned in such a way that its axis remains in a 10 mm diameter cylinder.

The leak rate has been measured as $Q < 3 \cdot 10^{-7}$ Pa·l·s^{-1}. A residual pressure of $< 10^{-4}$ Pa has been obtained. It must also be noted that there is no need to perform a bake-out of this tube.

4.5 The large valves

It is mandatory to have the possibility to close the whole volume of each tube (3.400 m^3) while venting the towers, so large valves are necessary at both tube ends. Such valves, with ~ 1 m aperture, are not available off-the-shelf for ultra-high vacuum. Their study and design was undertaken by Virgo; producing the following main requirements [19]:

- the valve must be a gate valve with a horizontal axis, made of stainless steel with a 1 m aperture diameter;
- the gate gasket must be in elastomer (Viton), 15 mm diameter cord;
- the leak rate must be $< 2 \cdot 10^{-7}$ Pa·l·s^{-1} for Helium, for the gate and $< 5 \cdot 10^{-5}$ Pa·l·s^{-1} for the gate body;
- the outgassing rates are $<5 \cdot 10^{-6}$ Pa·l·s^{-1} for H_2, for all stainless steel parts, and $< 4 \cdot 10^{-5}$ Pa·l·s^{-1} for the Viton gasket;
- the working differential pressure at the gate is 100,000 Pa in both directions, and 1000 Pa max. at opening or closing;
- the valve can be baked-out, when closed, at 150° for 100 hours, with its own heating device;
- the valve is fixed to the ground by means of 2 feet supports. The tightness must be kept when a 12 tonne stress due to atmospheric pressure is applied along the horizontal axis on either of its sides;
- the gaskets for the junctions of the valve to the tube links are made of stainless steel;
- the opening and closing times must be less than 60 seconds.

After a call for tender, the VAT-SARL firm (Verrières-le-Buisson, France) was retained, the valves being built at the VAT factory (Haag, Switzerland). A complete valve, 4200 mm high, 1610 mm wide and 600 mm thick, is shown in figure 26.

The gate is moved by a pneumatic actuator, separated from the vacuum by a bellows, and controlled by a double-effect electrovalve receiving air from a pressurised air bottle. The time required to open and close the valve have been respectively measured as 53 s and 30 s.

2012 JINST 7 P03012

Figure 26. A large valve.

The stainless steel pieces (except the piston bellows) have been air-fired at 400°C for 100 hours in the Virgo oven at SDMS. The Viton cord received a special thermal treatment (developed by Virgo) in three steps, in order to eliminate Fluorine, which could corrode the mirrors.

The heating device is made of two mobile insulated boxes on each side of the valve, which can roll and form a large closed box around the valve, with resistors and a fan to sweep the warm air,

2012 JINST 7 P03012

so heating the valve to 150° C by convection. This action has to be done at the same time as the arm tube bake-out or during the bake-out of the adjacent tower.

The four valves were all tested at VAT, and then again at Cascina following their installation in the halls. All of the requirements were met.

The complete control of the four valves, along with all the vacuum controls, is detailed below.

4.6 Tower and tube assembly

4.6.1 Survey

Knowledge of the positions of the vacuum system components is fundamental, as they support all of the interferometer optical elements. Tubes and towers have been positioned with respect to the precision reference bushes installed in tunnels and buildings, as described in the Infrastructure section. Taking into account that fabrication tolerances are ± 5 mm for tower bases and ± 2 mm for tube modules, main installation precision requirements are:

- misalignment between towers in the Central Building: ± 5 mm;
- misalignment between input and end towers of each 3 km arm: ± 5 mm (in addition to ± 5 mm measurement errors, due mainly to GPS);
- misalignment between adjacent tube supports: ± 2 mm;
- the centre of every arm tube cross-section must be inside a 100 mm diameter cylinder.

It has also been necessary to take into account the curvature of the Earth in order to have really straight arm tubes. Earth curvature also produces a 48 mrad angle between vertical directions over a 3 km distance; the same angle exists between the horizontal floors of the Central and Terminal buildings. Moreover, the West arm has an average inclination of about 1.5 m, over 3 km, since the West Terminal Building floor has an elevation that is lower by 1.5 m, with respect to the elevation of the ground level of the Central Building. All of these features have been accounted for in correctly positioning the tubes and towers.

4.6.2 Tower positions

The interferometer optical elements must be centred on the axes of the light beams, which are not straight lines, since they cross the beam splitter at a 45^{o} angle. Moreover, a few optical elements have a wedge angle between the two faces; this makes the outgoing beam not parallel to the entering one. On the other hand, each optical element, along with its support wire, is nominally aligned on the mechanical axis of the respective tower; the adjustment features of the inverted pendulum top stage and of the Separating Roof have been exploited to correct tower mechanical imperfections and positioning errors. As a consequence the towers are located, in the Central Building, according to the beam propagation path. In the Virgo reference system, where the horizontal orthogonal axes coincide with the interferometer arms, pointing approximately towards north and west:

- Injection and Power Recycling towers are translated by 20 mm westward;
- Beam splitter and Signal Recycling towers are translated by 20 mm northward;

Table 2. Horizontal coordinates and adjustment ranges of tower centres in the Virgo reference system.

Tower	West coordinate (mm)		North coordinate (mm)	
Input Bench	+20		-11000	
Power Recycling	+20		-6000	-600/+1000
Beam Splitter	0		+20	
North Input	0		+6400	±400
North End	0		+3006400	±1500
Detection Bench	-11000		0	
Signal Recycling	-6000		+20	
West Input	+5600	±400	0	
West End	+3005600	±1500	0	
Mode Cleaner	+144000		-11000	±2000

- the link tubes, where necessary, are connected off-centre to the towers, according to the given translations.

The towers are placed on the building floors, in correspondence with the holes giving access to the underlying payload installation gallery. The holes have been shaped in order to allow some adjustment and some asymmetry in the tower positions along the arm directions, as required by the interferometer locking scheme. The horizontal coordinates of tower centres in the Virgo reference system are listed in table 2, together with adjustment ranges.

It must be made clear that tower position adjustments are major changes, since they involve cutting link tubes and/or arm tubes and displacing very large objects, of the order of 20 tonnes each. The MC tower (the easiest one to move) was displaced in 2003 by 800 mm, requiring about two months of work. Very small displacements, of the order of millimetres, can be 'easily' performed, as long as they can be accommodated within the elasticity range of the link tube bellows.

Tower positions in the Central Building were measured at the very beginning, when only tower bases were in place and all lines of sight were free of obstacles. Subsequently, following oven and scaffolding installation, only very limited checks have been possible. There is evidence, however, that the building behaves like a rigid box. Moreover, the Central Building seems to be very stable with respect to the site external topographic references. The position of the only tower contained in each Terminal Building is, by definition, stable with respect to the stable floor plate, where it is clamped. The plate, in spite of its strong foundations, is steadily sinking at a relatively fast pace: about 20 mm/year at the West terminal and 10 mm/year at the North terminal. However, it is not a problem for the interferometer control system to follow such displacements by angularly steering the central mirrors by several μradians.

4.6.3 Tower installation

During the tower installation process, the most challenging operation was the introduction of the seven tower bases into the Central Building through the foreseen ad hoc trap-doors. It has to be considered that each base is a pyramidal object with a 4 m x 4 m square base and a weight of 17 tonnes. The tower bases, arriving from France on a special transportation chariot, were lifted by a

giant mobile crane and then lowered into the building through the trap-doors. Once on the floor, the bases, sliding on strong rollers, were pulled into position by capstans. The correct positioning, to within a few mm, was achieved in a few iterations, for the two horizontal coordinates. The vertical positioning was obtained with suitable shimming, which was also necessary as thermal insulation, during bake-out. A similar procedure was followed when installing the tower bases in the Terminal Buildings and in the Mode Cleaner Building, facilitated by having the external road at the same level of the inner floor.

The system to clamp towers to ground must be strong enough to stand unbalanced forces due to atmospheric pressure. The worse case is during standard operation, when the end towers and the arm tubes are in vacuum: in this situation the end tower base is subject to an 11 tonne force toward the Beam Splitter. On the other hand, during bake-out, the tower base must be free to thermally expand (by about 8 mm). The solution involved very strongly clamping to ground the oven/scaffolding structure containing the tower and leaving the tower square base plate free to expand on three sides, against suitable springs. The springs push the plate against the fourth side, serving as a fixed reference. The fixed side was chosen for the different towers according to the particular situation of atmospheric pressure imbalance.

With the tower bases in position, the oven/scaffolding structure was assembled and strongly clamped to ground. Then the systems of shims and springs mechanically connecting towers and ovens were installed.

The main successive operations were: external cleaning; pumping system installation and vacuum testing; bake-out oven testing; intermediate vacuum chamber installation and positioning; link tube installation; SA installation; payload installation; and final vacuum testing. All of these operations are described in the relevant sections of this paper.

4.6.4 Tube module preparation

Tube modules arrived by truck from the Belleli premises in Mantova, already fully equipped with thermo-couples and thermal insulation, and still sealed with tight end disks and breathing filters. They were first carefully vacuum cleaned, removing dust on the outer surface while passing through a transit tent, and then entered the assembly hall, at mid-length of each arm tunnel. The preservation of the cleanliness of the assembly hall and of the adjacent tunnel was considered very important, since the modules had to be opened in order to be welded to one another in this environment. In the hall, serving also as buffer storage for up to 10 modules, the plastic envelops protecting the lips and the end disks were dismounted. Then, with the help of a bridge crane, each module was loaded on to an electric train and sent into the appropriate half tunnel.

4.6.5 Tube support installation

Tube supports were installed in the tunnel as soon as the tunnel sections were available and equipped with surveyed reference bushes, well before the delivery of the modules. The supports were bolted to the tunnel floor and the four upper plates of each support were individually positioned to ± 2 mm with respect to the local reference bush, ready to accept the module feet.

2012 JINST 7 P03012

4.6.6 Tube module positioning and welding

At installation, the two front feet of the new module were posed on the tube support already holding the two rear feet of the previously installed module. The bellows at the interface between the two modules was kept compressed with an appropriate tool, in order to leave room for the still present end disks, protecting the lips to be welded. The following operations were performed in sequence:

- fine adjust the new module position with respect to the tunnel reference bush, using the precision markers drawn on the modules, at the factory;
- slide over the welding zone a movable clean tent, equipped with filtered air supply, enclosing about 6 m of tube, including the support, and wait for the establishment of a clean environment;
- clean with solvents all the surfaces inside the tent;
- start blowing filtered air into the already installed tube and into the new module;
- as soon as the overpressure is sufficient, slide off the two adjacent end disks;
- clean the lips;
- release the compressed bellows, bringing into contact the pair of lips to be welded;
- fine adjust the position of the new module feet, in order to obtain a perfect lip correspondence, to better than 0.5 mm;
- install around the tube the welding robot and bring into the tent the PC controller;
- let the robot chariot perform a test full turn around the lips;
- start the welding turn and check continuously the electric parameters and the melted metal shape on the screen;
- inspect carefully the whole weld;
- dismount the robot and slide away the tent.

A visually un-satisfactory weld occurred on only a few occasions, in general due to inaccuracy of the position of the torch with respect to the lip edges. The defective weld portions were corrected by a second torch passage.

The initial programme foresaw the assembly of 300 m-long tube sections (20 modules, 19 welds), with welded end caps, to be He leak-tested separately. Following the absence of leaks in the first two sections, we decided to build longer sections, in order to speed up the assembly. The last assembled section was as long as 1.5 km, i.e. the whole last half of the West arm.

Out of more than 400 circular welds, only two were leaking and required to be cut by the nibbling robot and re-welded on to the fresh lip edges.

The cutting head of the robot was also used to cut-out the welded end caps of the tested sections, before joining them together with appropriate 1 m long bellows sections.

2012 JINST 7 P03012

4.6.7 Tube leak detection

Tube leak detection was performed with a classical procedure, made exceptional by the enormous volume of the system, in particular when a whole 3 km tube was tested for the section-joining welds. Waiting times larger than 1 hour were necessary.

Every circular weld to be tested was enclosed in a band of silicon rubber (standing bake-out temperature) glued to the tube wall, on the two sides of the welded lips. Having connected an He leak detector to the pumping port, gaseous He was fed into the volume under the rubber band surrounding the lip weld. According to the design specification a maximum acceptable leak rate of $3\ 10^{-11}$ Pa m^3 / s ($3\ 10^{-10}$ mbar l / s) was used.

A test bake-out was performed on all of the tube sections, but the leak test was not repeated after bake-out, since the reached ultimate pressure and the residual gas composition witnessed the absence of harmful leaks.

4.6.8 Arm tube geometrical stability

Ground subsidence has to be continuously corrected for two main reasons.

First, because a support sinking much more than the support at the other end of the same module may produce an angle between two modules, hence stress in the lip welds. The maximum acceptable angle has been evaluated to be about 0.7 mrad, corresponding to one support sinking by 10 mm, between two completely stable supports. Before reaching this limit, the relative displacements of the modules can be safely accommodated within the elasticity of the bellows and lips and in the tolerances of bolted connections.

Second, because the tube axis must be straight within $\pm$50 mm. Beyond this limit tube walls and baffles start scraping the beam tails. This limit was retained, even if it was initially evaluated for the case of three interferometer beams in the tube.

Tube supports have been conceived in order to allow vertical tube realignment to recover tunnel subsidence. This is easily and safely achieved, using hydraulic jacks and adjustable safety shims, keeping the tube under vacuum. It turned out that there are zones, involving up to 20 supports, with a maximum sinkage of 20 mm/year. In this case two realignment operations per year are necessary.

Particular care has to be paid at the far arm ends, since the last tunnel slabs sink more rapidly than the terminal buildings, producing a share stress on the tube. As, at these points, the tube is not equipped with bellows, the maximum acceptable relative sinkage is only 2 mm, requiring more frequent adjustments. These actions have been made very easy and quick, by the replacement of the standard legs with mechanical jacks, on the two last tube supports. Very high precision hydraulic levels are being tested to monitor subsidence in these zones.

4.7 The pumping system

4.7.1 Pumping system function and concept

The pumping system includes all the devices necessary to produce and monitor the vacuum environment for the proper functioning of the Virgo interferometer. The vacuum enclosure is composed of several volumes, which are separable by gate valves, in order to make their vacuum status independent from one another; for the same reason, each separable volume has its own pumping and

control system. There are 12 separable volumes, but the two 3 km arm tubes have 6 pumping stations each, to make their pressure profile sufficiently uniform, over the huge length. Moreover, the six long towers each have two main compartments, with different vacuum regimes. This gives a total of 22 pumping stations, distributed over 6 kilometres; to make control and maintenance easier, standardisation of components has been enforced as much as possible.

All pumps and gauges are connected to the vacuum tank through valves. In this way, in case of failure, the defective equipment is separated (automatically, by the remote control system) from the vacuum tank, without interrupting interferometer operation. Given the very high conductance existing between all parts of the vacuum tank, shutting down one pumping group for a limited time does not qualitatively change the vacuum level.

The valves separating the arm tubes from the rest of the vacuum system have a 1 m diameter aperture and have been custom-designed and built. The 'Large Valves' have already been described in a dedicated section (4.5).

All the other valves are standard, with aperture diameters up to 250 mm; shutter gaskets are in Viton, except valves normally closed against atmosphere, where metal gaskets are used.

The whole pumping system, valves included, is remotely controlled through one control/command crate per pumping station, linked to the Slow Monitoring network of Virgo. The control can be operated from the main workstations in the Control Room, from a portable PC, connected locally to the network or, exceptionally, in local manual mode, in case of network failure. Any action on pumps, valves and gauges is subject to hard-wired and software safety conditions, in order to make wrong or dangerous operations impossible.

The main vacuum requirements, described in more detail at the beginning of the vacuum chapter, are 10^{-7} Pa (10^{-9} mbar), in the UHV volumes, and 10^{-4} Pa (10^{-6} mbar), in the HV volumes.

Since the UHV regime is reached through a 150 ^{o}C bake-out, all of the connected valves, pumps and gauges have to be able to withstand the same treatment and have been equipped with appropriate heaters.

4.7.2 Cleanliness

In addition to maximum pressure requirements, there is also the cleanliness issue: hydrocarbon vapours below 10^{-11} Pa (10^{-13} mbar) and absence of dust. All lubricants are forbidden and great care is put into avoiding dust pollution. For this reason, most of the primary pumps are of the dry type ('scroll') and all the large turbomolecular pumps have magnetic not lubricated bearings. This solution is also fundamental to the reduction of vibrations, which are obviously very problematic for an apparatus designed to detect gravitational waves.

For the sake of hydrocarbon and dust cleanliness, all electromechanical components to be installed in vacuum have been carefully selected and tested; appropriate solvents have been used and heat treatments under vacuum have been developed.

4.7.3 Tower pumping

As previously explained, in the long towers there is a bottom UHV compartment, where the suspended optical element is located, and a top HV compartment, in which the anti-seismic suspension is found. This device contains several electromechanical components, magnets, motors and

2012 JINST 7 P03012

hundreds of metres of cables, summing up to a surface area of plastic material of about 100 m^2. Between the two compartments there is the Separating Roof, described in a previous section (4.2.4), confining an intermediate volume, to be pumped independently, in order to limit the gas flow from the HV compartment to the UHV compartment.

The top and the bottom compartments are pumped by two similar groups, each one made of a 1600 l/s magnetic bearing hybrid turbomolecular pump and of a 25 m^3/h scroll pump, serving both as roughing pump and backing pump. It has been calculated and experimentally verified that during rough pumping the pressure difference across the Kapton windows of the Separating Roof is correctly limited and that the air flow is not strong enough to move accumulated dust particles.

HV regime is reached in two steps: 'rough pumping', from atmosphere to 10 Pa with the scroll pumps, and 'intermediate pumping', down to 10^{-4} Pa, with the turbomolecular pumps. To reach UHV in the bottom compartment a third step is necessary, continuing the turbomolecular pumping action, with the addition of bake-out. After cooling down, when the pressure is below 10^{-6} Pa, the 'permanent pumping' is started in order to reach 10^{-7} Pa or less. This last step is performed by opening the valves in front of the Ti evaporation pumps, each consisting of two 250 l tanks equipped with Ti evaporators. The effective pumping speed, about 3000 l/s, is determined by the conductance of the connection pipes; the saturation of a fresh Ti layer is reached after about one year of operation. One of the tanks gives access to a 200 l/s Ion pump, necessary to pump noble gases, not captured by titanium. The Ti evaporation pump is connected to the tower lower compartment through one of the link tubes.

During evacuation, the intermediate vacuum chamber is pumped by the upper and lower pumping groups, through the conductance pipes. When the bake-out is nearly finished, in order to allow to reach 10^{-7} Pa in the lower compartment, an Ion pump, connected to the intermediate chamber by a bellows, is put into operation, to achieve the wanted gas flow reduction between top and bottom compartments.

In standard operation conditions, the permanent pumping is constituted by the turbo/scroll group operating on the top compartment, the ion pump operating on the intermediate chamber, and the Ti evaporator/ion pump operating on the mirror compartment.

Thanks to the very low outgassing of the tower walls and of the SA, it has been demonstrated that it is possible to operate the scroll pump backing the turbo only five minutes per week. This is very important, as it enables the elimination of a possible source of mechanical noise.

In a system where such huge volumes have to be evacuated, the duration of all of these operations has to be carefully considered. The sizes of the pumps chosen to meet the pressure requirements give relatively convenient evacuation times.

For both long and short towers:

- rough pumping: about 16 hours;
- intermediate pumping: a few hours to reach 10^{-3} Pa.

This means that a tower can be evacuated overnight.

The bake-out time, at the chosen temperature, is much larger: about one week, including cooling down. Due to this fact, Virgo is presently operating in pre-bake-out HV regime. This is necessary because, during commissioning, the towers have to be opened rather frequently and

2012 JINST 7 P03012

would require a bake-out after any evacuation, in order to reach UHV regime. This is not a problem, since the intermediate pumping system already allows the reaching of a pressure close to 10^{-5} Pa, in the mirror compartments. This pressure value is absolutely acceptable from the point of view of noise induced by vacuum fluctuation; a lower value would be required if all other noises could be pushed well below the present design sensitivity. A few bake-out runs have been performed for test purposes.

Tower venting time has been set at about 14 hours (approximately one night). This is due to the need to avoid moving dust particles and to limit to a few hundred Pa the maximum pressure difference across the Kapton windows of the Separating Roof. The impedance created by dust filters and air dryers on the inlet line helps to this end. Given the large volume, venting air comes from the atmosphere; hence it has to be dried to avoid potentially dangerous water condensation.

All that has been stated for the long towers also holds for the short, limited to the HV regime; since short towers have one single compartment, no bake-out is performed and there are no Ti evaporation pumps.

The Mode Cleaner tube, 142 m long, is cut by a gate valve close to the IB tower. Hence the MC tube is evacuated together with the MC tower.

Also the Signal Recycling tower has one single compartment, but being part of the UHV volume, it is equipped with a Ti pump.

4.7.4 Tube pumping

Each 3 km arm tube is equipped with six pumping stations, at 600 m distance from one another, which are very similar to the tower lower compartment pumping stations, including a 1600 l/s magnetic bearing hybrid turbomolecular pump backed by a 25 m^3/h scroll pump. There is also a single volume Ti evaporator with a 200 l/s ion pump. The turbo/scroll group is used as intermediate pumping and during bake-out, the Ti/ion group is used as totally vibration-free permanent pumping, in the regime 10^{-6} Pa to 10^{-8} Pa.

Due to the huge volume, the tubes cannot be rough pumped by the scroll pumps. In fact, it is necessary to employ a group able to pump against atmospheric pressure for more than one day. This heavy duty is performed, on one tube at a time, by a large group comprised of a 1000 m^3/h Roots pump, backed by a 250 m^3/h oil-free pump. The roughing group ultimate pressure is 10^{-1} Pa, but intermediate pumping is started at 10 Pa, in order to avoid any backstreaming.

Tube pumping times are much longer than those of the towers:

- rough pumping: about four days;
- intermediate pumping: about one day.

The chosen bake-out time is two weeks, to compensate for possible temperature disuniformities in such a long system.

Such long times have no drawback, since the large valves allow the maintainenance of the tubes in vacuum, when opening the towers. On the other hand, since the tubes do not contain any equipment, there is no need to vent them. In fact, the arm tubes have been evacuated only once, in 2003, and they have not been vented, since then.

2012 JINST 7 P03012

Tube venting time is of the order of a few days; venting is performed, when needed, using filtered atmospheric air.

As is the case with the towers, also the tubes will only be definitively baked-out at the end of the commissioning phase, when venting a tower will very seldom occur.

4.7.5 Control instruments

Most control instruments and vacuum gauges are installed on dedicated ports ('Instrumentation Bottles'), connected through valves to the vacuum system. In this way, in case of failure or maintenance, they can be accessed without venting the large volumes. In general, there is one instrumentation bottle per pumping group.

Different types of pressure gauges have been chosen, according to the required performance and to the pressure range: Pirani; Capacitance; Inverse Magnetron; and Bayard-Alpert. On almost any instrumentation bottle there is also an RGA. This instrument has proven very useful in monitoring the vacuum quality, monitoring hydrocarbon contamination and detecting leaks.

High sensitivity differential gauges have been installed across the Large Valves, to prevent operations with pressure differences higher than 2000 Pa.

All the instruments are read and controlled remotely, as described in detail in the following sections.

4.8 Vacuum control system

The vacuum control system [20] has been designed in such a way that the main vacuum elements of the Tube (Tu), Tower (To), Link Valve (LV) and Large Gate Valve (LGV) stations can be controlled and monitored remotely. This includes all basic actions needed to reach the vacuum level required in the different parts of the vacuum chamber and the continuous acquisition of all slow control channels providing information about the state of the vacuum itself (pressure values, pump activity, valve status) as well as the vacuum control elements.

Moreover, the vacuum control system includes safety protections for all actions considered hazardous for personnel or dangerous for equipment. The philosophy of the safety protection architecture is to introduce redundancy for all protections against hazardous commands: at least two completely independent safety conditions are required, one of them being hardwired. Hardware safeties are implemented in every bench (Tu, To and Valves), and software safeties are implemented in a local server for Tu and To, and in the top Supervisor (VTTVSU) for the valves. The software safeties enable the combination of information from the different parts of the vacuum chamber. For instance, the large valves cannot be opened if the two capacitance gauges installed on each side of them read a pressure difference greater than 500 Pa (5 mbar). The safety protection can be conditional or active: a conditional safety forbids, for example, to act on a pump if a condition is not fulfilled. An active safety permanently monitors the activity of a pump and executes a specified action when a given condition is detected [21].

4.8.1 Vacuum control hardware

There are 6 pumping benches for each tube and 10 pumping benches for the towers, among them are 7 benches for UHV (long) towers and 3 for HV (short) towers, as shown in figure 27.

2012 JINST 7 P03012

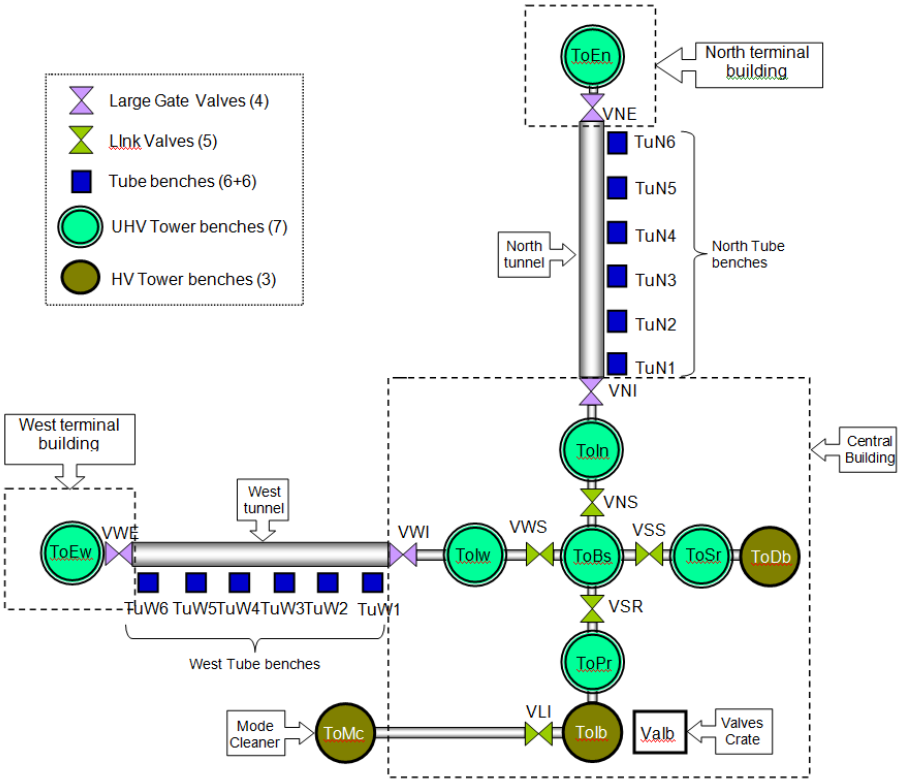

Figure 27. Schematics of the vacuum control system.

A pumping bench is typically in charge of controlling the pumps and valves and reading gauge pressures. Power to the benches is delivered by dedicated UPS lines. Each bench is composed of an electrical part (power breakers, relays, power supplies) and an electronic crate managed by a CPU board (Motorola 68040 processor) on which runs the OS9 Operating System. The tube and tower pumping benches are based on the same electronic architecture, while the electrical part is very specific as a tower vacuum bench is much more complex. In addition to this, the racks contain specific controller boards associated to the pumps, valves and gauges. At the start-up of the electronic crate, as a first step, a boot programme is uploaded from a dedicated server (*os9boot*) located in the Data Acquisition (DAQ) room, in the Central Building, which, in a second step loads into the CPU memory the specific server programme associated with Tube or Tower functionalities, which then runs permanently.

The valves are of two kinds: there are 5 link valves (LV) located between the towers and 4 large gate valves (LGV) in order to isolate the tubes from the rest of the vacuum chamber. These 9 valves are controlled by one crate, hosted in the Injection Bench tower pumping rack, linked through a local network FIP to seven control boxes (one for each valve except for the 2 large gate valves located in the Central Building, which share a control box with a link valve). The valve crate implements the same CPU as the tower and tube crates.

2012 JINST 7 P03012

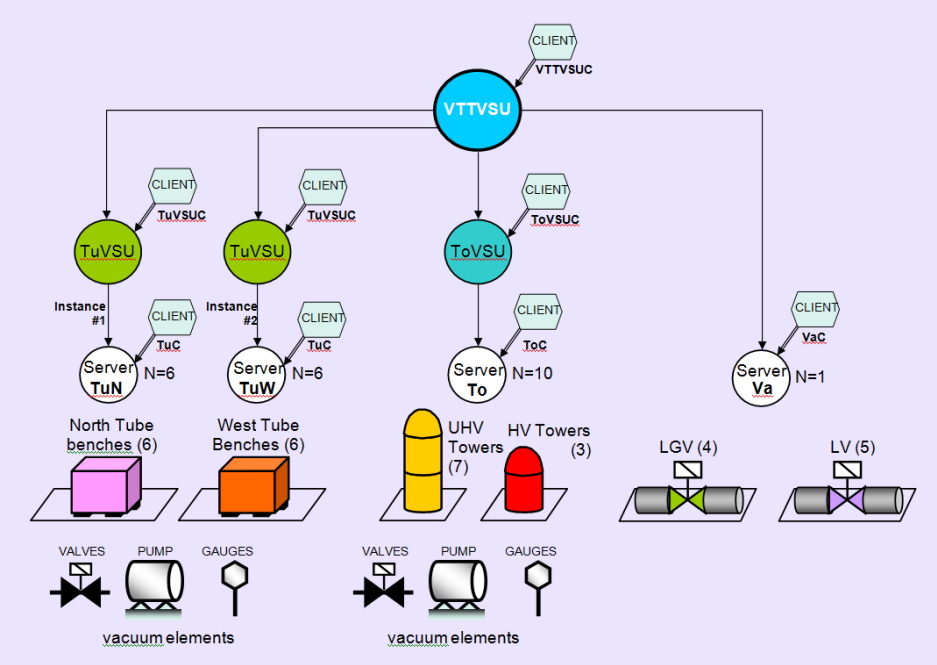

Figure 28. Schematics of the vacuum control software.

4.8.2 Vacuum control software

The vacuum control software is based on a client-server architecture. The servers running on the OS9 CPU are unique: there are 12 tube servers, 10 tower servers and 1 server for all the valves. The vacuum control software has been developed in the framework of the Supervisor (see the section on Interferometer Control). It therefore benefits from this architecture. For instance, it makes use of the mastership communication protocol implemented in the Supervisor in order to arbitrate which software element is allowed to take the mastership on the control of a pumping bench at a given moment: only one client (a graphical user interface) can send orders to a server to perform a specified action, such as opening a valve or turning off a pump. Until the mastership has been released, no other client can obtain it.

The architecture of the vacuum supervisor is hierarchical, as shown in figure 28: the tube servers are controlled by two instances of a supervisor (TuVSU) controlling the 6 tube servers of each arm.

This choice was driven by the necessity to operate each arm independently of the other. The 10 towers are controlled by another supervisor (ToVSU). These 3 supervisor servers are then run through a top supervisor VTTVSU (VTT for Valve Tube Tower) which directly controls the valve server. The different VSU (ToVSU, TuVSU and VTTVSU) have specific user interface clients displaying some information about the status of the different vacuum control elements (pressure, gauge status, valves status, server status) which are typically refreshed each second. Alarms are also gathered and displayed in real time on the user interface. The vacuum VSU are also in charge of forwarding the alarms sent by the servers to the ErrorLogger daemon, which collects error messages coming from all servers running in Virgo.

2012 JINST 7 P03012

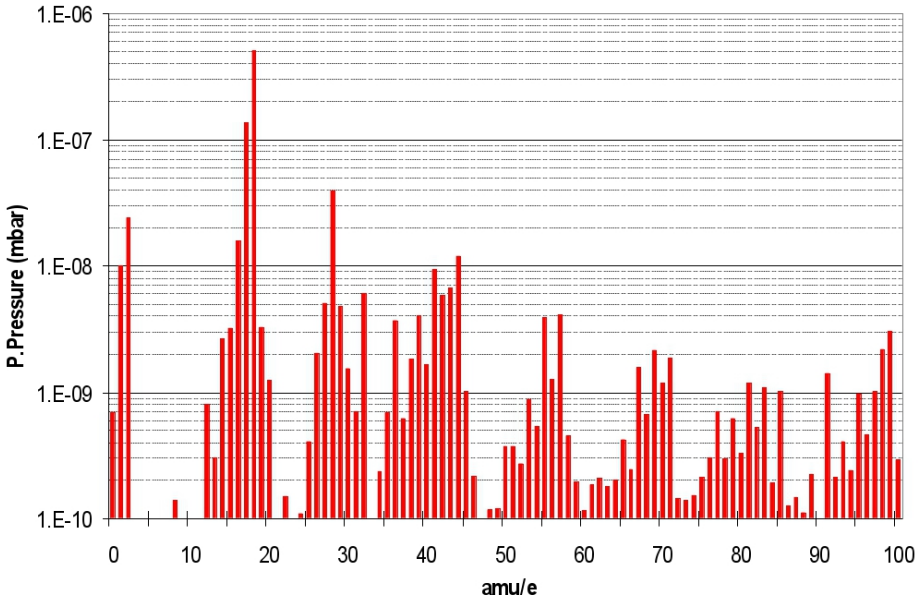

Figure 29. Residual gas spectrum for the upper part of a long tower.

The software safety protections for the tower and tube actions have been implemented in each tube and tower server, while they have been implemented in the VTT server itself for the valves because it requires the knowledge of some information (mainly pressure) provided by the tower servers.

Finally, the OS9 servers send slow monitoring information to the DAQ system (slow frame builders - Fbs) every 10 seconds in order to acquire vacuum data, which might be important in understanding the apparatus at a given time.

4.9 Vacuum system performance

As explained, at present, in the commissioning phase, the vacuum system is kept in pre-bake-out HV regime.

Most towers have been evacuated and vented several times. Each long or short tower is pumped by one scroll-turbo group and the backing scroll pump is only put into operation 5 minutes per week, when a consistent quantity of gas (up to a pressure of 500 Pa) has been accumulated by the turbo. A typical residual gas spectrum for the upper compartment of a long tower is shown in figure 29. Short towers behave in the same way as the upper compartments of long towers.

Bake-out of a tower lower compartment has only been performed for test purposes, so far, on towers fully equipped with Superattenuators. The residual gas spectra of a lower tower compartment, before and after bake-out are shown in figure 30.

The 3 km arm tubes have been evacuated only once, in 2003, and have not been vented since. As for the towers, the final bake-out has not yet been performed. At present, they can be kept well below 10^{-4} Pa, with only one scroll-turbo pumping group per tube, thanks to the excellent outgassing and cleaning procedures, adopted during fabrication and installation. In these conditions, the pressure profile along a tube varies between $0.4\text{x}10^{-4}$ Pa and $0.8\text{x}10^{-4}$ Pa.

2012 JINST 7 P03012

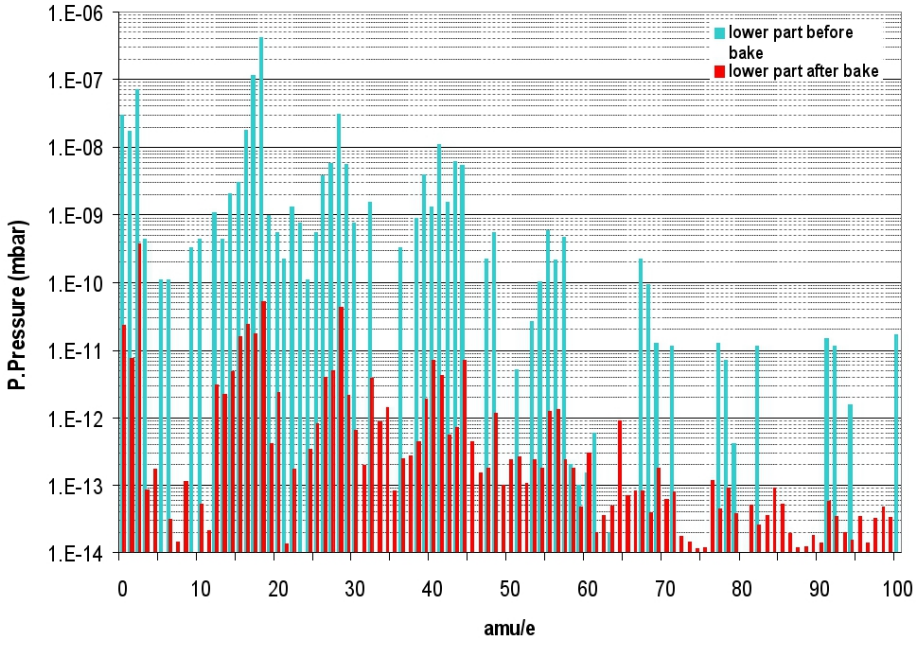

Figure 30. Residual gas spectrum for the lower compartment of a long tower, before and after bake-out.

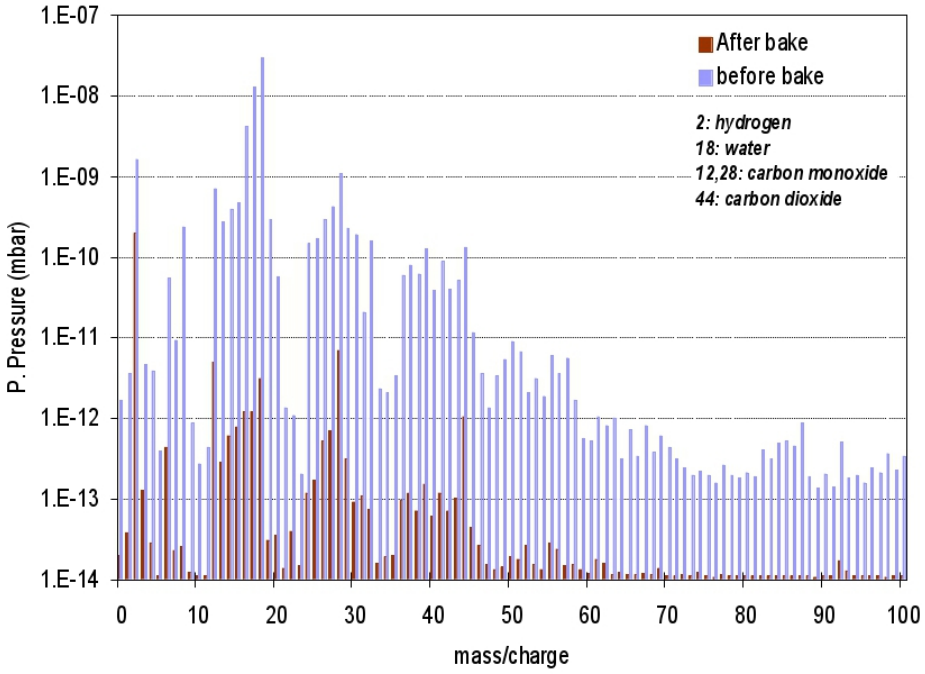

Figure 31. Residual gas spectrum for a 3 km arm tube, before and after bake-out.

One full tube has been baked-out for test purposes, achieving an average end pressure of $0.5 \text{x} 10^{-7}$ Pa. The residual gas spectra of a tube before and after bake-out are shown in figure 31.

5 Scattered light noise

5.1 Introduction

The sensitivity of the interferometer can be limited by noise due to scattered light propagation and recombination in the interferometer beams.

The total light scattered by the mirrors is of the order of 100 ppm, mainly due to mirror 'flatness' (i.e. large-scale polishing defects), resulting in diffraction at low angles (or in low order modes), and thus in a loss of power and contrast. Mirror roughness (small-scale defects) results in light scattered at 'large' angles ($> 10^{-4}$ rad) at the level of some ppm, and is responsible for the 'scattered light' noise considered here.

This scattered light propagates in the vacuum chamber, and, following reflections and diffusion from the walls, may recombine with the main beam, with a different phase, due to the different optical path. The vibrations of the chamber make the phase shift time dependent, resulting in noise in the phase of the interference signal.

5.2 Noise evaluation

A careful noise evaluation of possible 'killing' baffle systems was performed; the initial ideas of this evaluation having been formulated by K.Thorne [22], while we also developed a statistical approach to scattered light propagation [23].

The information on trajectories obtained from a particle analysis was supplemented by additional information regarding phase modulation coming from the wave picture. Scattered light was represented by optical rays, extracting information about the noise from the statistics relating to photons received by the mirrors. Subsequently, the options for a baffle system were evaluated with a Monte Carlo simulation. Coherent effects, such as diffraction or back reflection by the baffle edges, were computed analytically [24].

As a first result, the noise in a plain tube was evaluated to be $h(f) \sim 10^{-23}/\sqrt{\text{Hz}}$ at 10 Hz. Being of the same order as the foreseen sensitivity, it became obvious that a system of baffles would have to be designed and implemented in the arm tubes. The simulation was confirmed by analytical calculations.

5.3 Baffles in the arm tubes

As it had been decided to allow for the possibility of having three interferometers side-by-side inside the vacuum tank, requiring a free aperture of 1 m, a space in the tube cylinder between ϕ 1000 mm and ϕ 1200 mm was left available for baffles.

The best performing material was absorbing black glass for welding protection with an anti-reflective coating (Schott-Desag Athermal 11, thickness 3.2 mm, transmission $\sim 10^{-7}$, reflection between 3.10^{-3} and 10^{-2} depending on incidence angle, and negligible diffusion, at $\lambda = 1064$ nm). A satisfactory glass prototype was developed, with the shape of a quasi perfect light trap, but finally these baffles were found to be too fragile for installation in the 3 km tubes. The black glass was used only in the central area and in the links between valves and towers. The baffles in the 3 km arm tubes were built in stainless steel.

2012 JINST 7 P03012

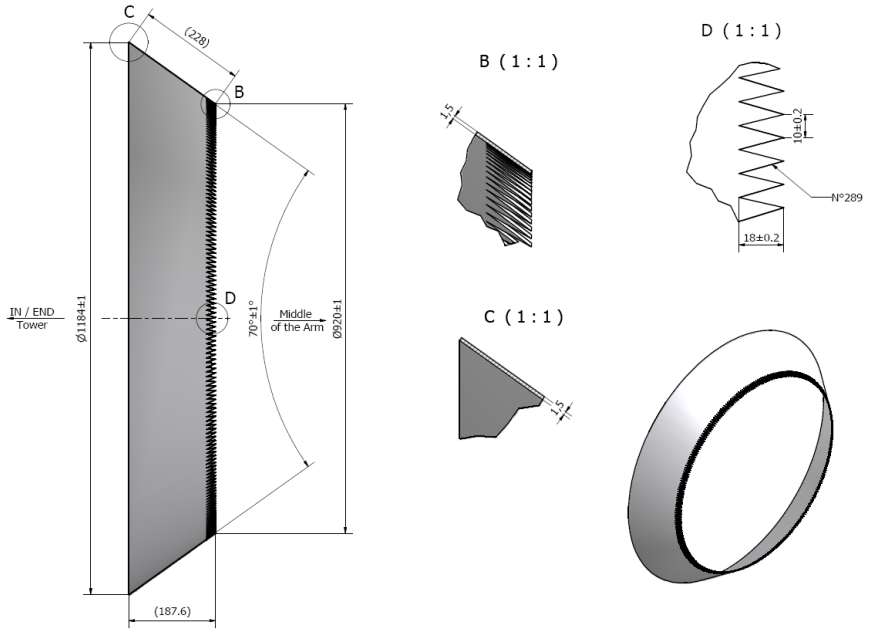

Figure 32. Shape of a conic metal baffle with serrations at the inner edge.

5.3.1 Steel baffles

Stainless steel baffles with a conic shape are more efficient than simple flat rings, perpendicular to the tube axis. Stainless steel has a small reflection coefficient, but a large reflectance distribution function (BRDF), and absorption is close to 50%. The BRDF was measured for the sandblasted stainless steel used for the baffles, and was introduced into the Monte Carlo computation.

Finally, steel baffles have the shape of a truncated cone with an external diameter of 1184 mm, internal diameter of 990 mm and an angle of 55° with respect to the cone axis. The steel thickness is 1.5 mm. The internal edge is serrated in order to suppress diffraction (figure 32).

The conceptual requirement was to completely cover the solid angle as seen by the beam spot at each Fabry-Perot mirror. A simple calculation gives the distance *Zn* of the nth baffle to the mirror through the law:

$$Zn = Z1\left(\alpha\frac{r}{r-h}\right)^{n-1}$$

where r is the tube radius, h is the baffle height, and α is a coefficient to obtain some coverage between the shadows of each baffle. The total number of steel baffles is 160 per arm. These were produced by several firms near Paris and soldered inside the modules at CNIM before the washing, by means of 4 small lugs on their external part. They were only installed in the tube part in between the two large valves, which is ideally supposed to remain untouched for the life-time of Virgo.

5.3.2 Glass baffles

In the links joining the large valves and the towers, large glass baffles have been installed to cover all discontinuities, such as valve or tower flanges, bellows, etc. Glass baffles are allowed in this area, since they can be relatively easily replaced, in case of breakage. All of these link baffles are

2012 JINST 7 P03012

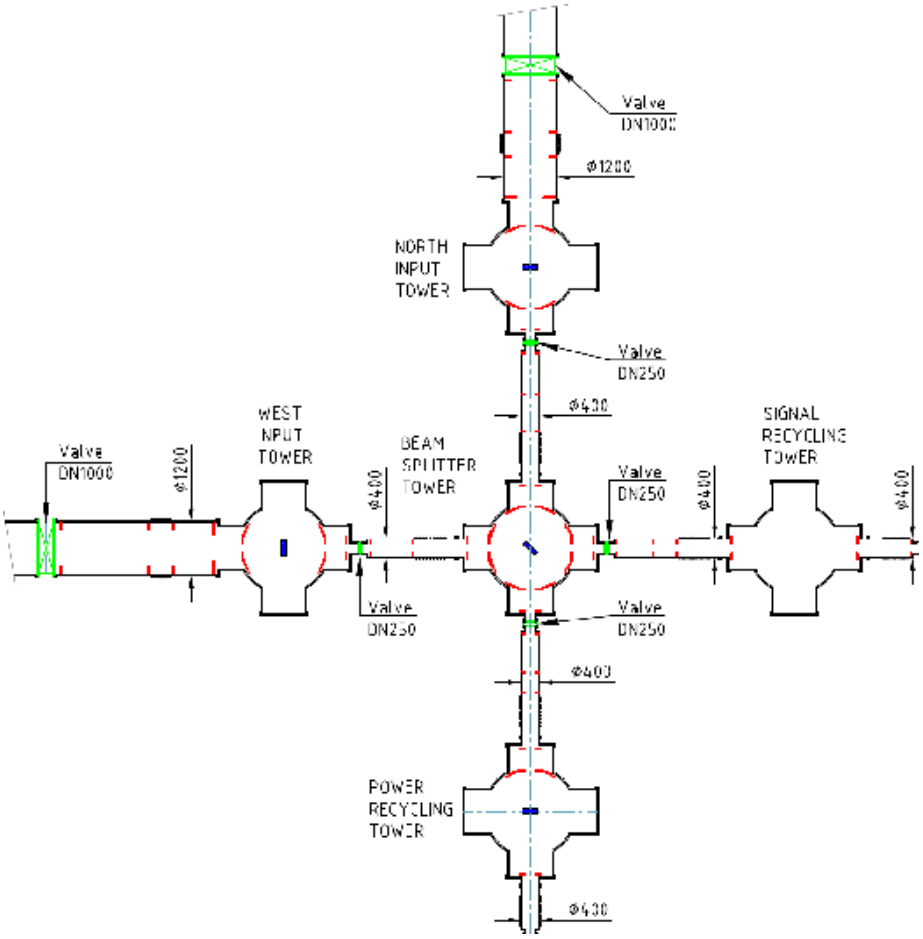

Figure 33. Scheme of the central area and location of glass baffles, represented in red.

1160 mm diameter disks made of black glass, with a large central hole. Inside each link are 4 rings, one in front of each flange (tower and valve), and one on each side of the bellows. The inner hole diameter ranges from 600 mm to 850 mm according to the baffle location and the laser beam size. Their sizes and positions are displayed in figure 33.

With these stainless steel and glass baffles, the evaluated noise in the arm tubes is $h = 6.5\ 10^{-25}/\sqrt{\text{Hz}}$ at 10 Hz.

5.4 Baffles in the Central Area

The central area comprises 5 towers, their mirrors (figure 34) and the 4 junction links (small links). To evaluate the noise with a Monte Carlo simulation, all possible processes for re-scattering of photons off all elements were listed. The geometry of the tanks, links and mirrors, and the optical properties of all surfaces, were taken into account [25]. Photons coming from the main beam were generated and their scattering off each element was followed in all possible paths until a possible recombination. The total noise is the incoherent sum of the contributions of all scattered light generation processes through all recombination processes.

The resulting actions for reducing the noise are given below, together with the final result.

2012 JINST 7 P03012

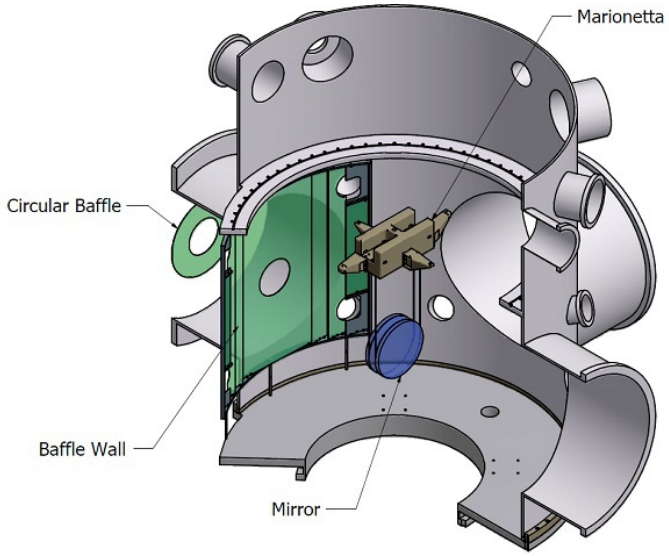

Figure 34. Glass baffles covering the tower inner wall around the laser beam path.

For the small links, the steel tube diameter must not be less than 400 mm, to avoid the effects of grazing diffusion on its walls. Black glass disks with an external diameter of 380 mm, and a central hole, with a diameter of 230 mm, corresponding to ± 5 beam waists, were inserted inside the links. Their position was chosen in order to hide the 250 mm diameter valves and the bellows waves (figure 33). With these baffles (14 altogether) the noise coming from the small links is negligible.

Among the towers the most delicate situations can be found around the two input mirrors and the Beam Splitter, due to the complex geometry and the high light intensity. In fact, all possible diffusion or reflection areas in the tanks have to be covered with black glass, with ± 5 waist holes for beam passage. In order to be as close as possible to the tower cylinder, the baffles inside the towers at the junction of the big links have the shape shown in green in figure 34: in the centre a big plate 1150 $\times$ 600 mm^2 with a hole having a radius equal to 5 times the beam waist, and on each side three 1150 $\times$ 120 mm^2 plates installed as shown in the side view in figure 34. The Beam Splitter tower has 4 such devices, covering most of the cylinder (figures 33 and 35), the input towers have 2, the end towers have one and the recycling tower also has one.

The resulting noise from the central area is reduced to $\sim$ some $10^{-25}/\sqrt{\text{Hz}}$.

2012 JINST 7 P03012

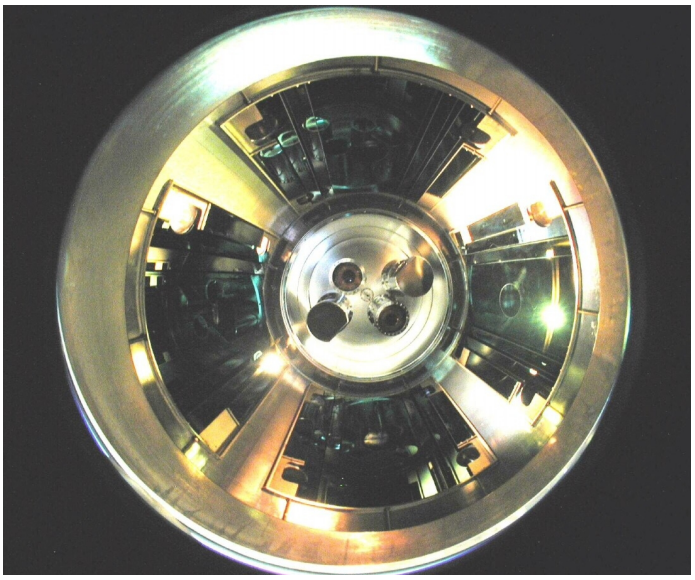

Figure 35. Picture of the 4 glass baffle sets inside the Beam Splitter tower, taken from below with a fish-eye lens before mirror installation.

2012 JINST 7 P03012

6 Mirror suspension

6.1 Seismic isolation

Seismic motion, always present on the ground, even in the absence of earthquakes, is a very powerful noise source for gravitational wave detectors. The typical linear spectral density of seismic displacement in all three space coordinates can be approximated, in the frequency range of Virgo, by:

$$\tilde{x} \simeq \frac{\alpha}{f^2}$$

measured in m $Hz^{-1/2}$, as a function of the frequency f and of the parameter α, which ranges typically between 10^{-6} and 10^{-9} m $Hz^{3/2}$, depending on the geographical location.

From the previous equation, it turns out that the seismic motion reaching the test masses (the mirrors) has to be reduced by at least a factor 10^{12}, in order to have an acceptable sensitivity in the 10 Hz region.

High-performing devices have been developed in order to suspend the Virgo mirrors, strongly reducing the seismic-induced motion; those devices have been called: 'Superattenuators' (SA).

6.2 Superattenuator working principle

The SA design is based on a simple idea. If the suspension point of a simple pendulum moves in the horizontal plane at a frequency f higher than the pendulum normal mode (f_0), the oscillation is transmitted to the suspended mass with attenuation proportional to $(f_0/f)^2$. The wire length l connecting the suspension point to the mass sets the pendulum normal mode frequency according to the equation:

$$f_0 = \frac{1}{2\pi}\sqrt{\frac{g}{l}}.$$

Therefore, a device suspending a mirror based on the working principle of a simple pendulum, represents a good isolation system against horizontal seismic noise at frequency $f \gg f_0$.

A better attenuation performance is obtainable with a pendulum chain, where each pendulum is attached to the mass of the one above, forming an n-stage pendulum (figure 36).

With this system an oscillation at a frequency f higher than the frequencies of the chain normal modes ($f > f_n > \ldots > f_1$), is transmitted to the suspended mass with attenuation proportional to f^{-2n}. In particular, the ratio between the linear spectral density of the last mass displacement (the optical component) and the linear spectral density of the suspension point displacement (where the excitation is applied), decreases as A/f^{2n} where $A = f_1^2 \cdot f_2^2 \cdot \ldots f_n^2$ and n is the number of stages. In this way a very large attenuation of the seismic noise horizontal component can be obtained in the frequency region above the highest pendulum resonance frequency, simply by increasing the number of stages.

Unfortunately, the pendulum wires supporting the end mirrors, at the extremities of the 3 km arms of Virgo, are misaligned with respect to one another by about $3*10^{-4}$ rad, because of the different directions of the plumb line on the curved surface of the Earth. This means that they are not perfectly perpendicular to the laser beam, hence vertical seismic motion introduces significant changes in the distance between the mirrors. In addition, any vertical vibration will be partially transmitted to the interferometer horizontal axis (laser beam direction) because of unavoidable

2012 JINST 7 P03012

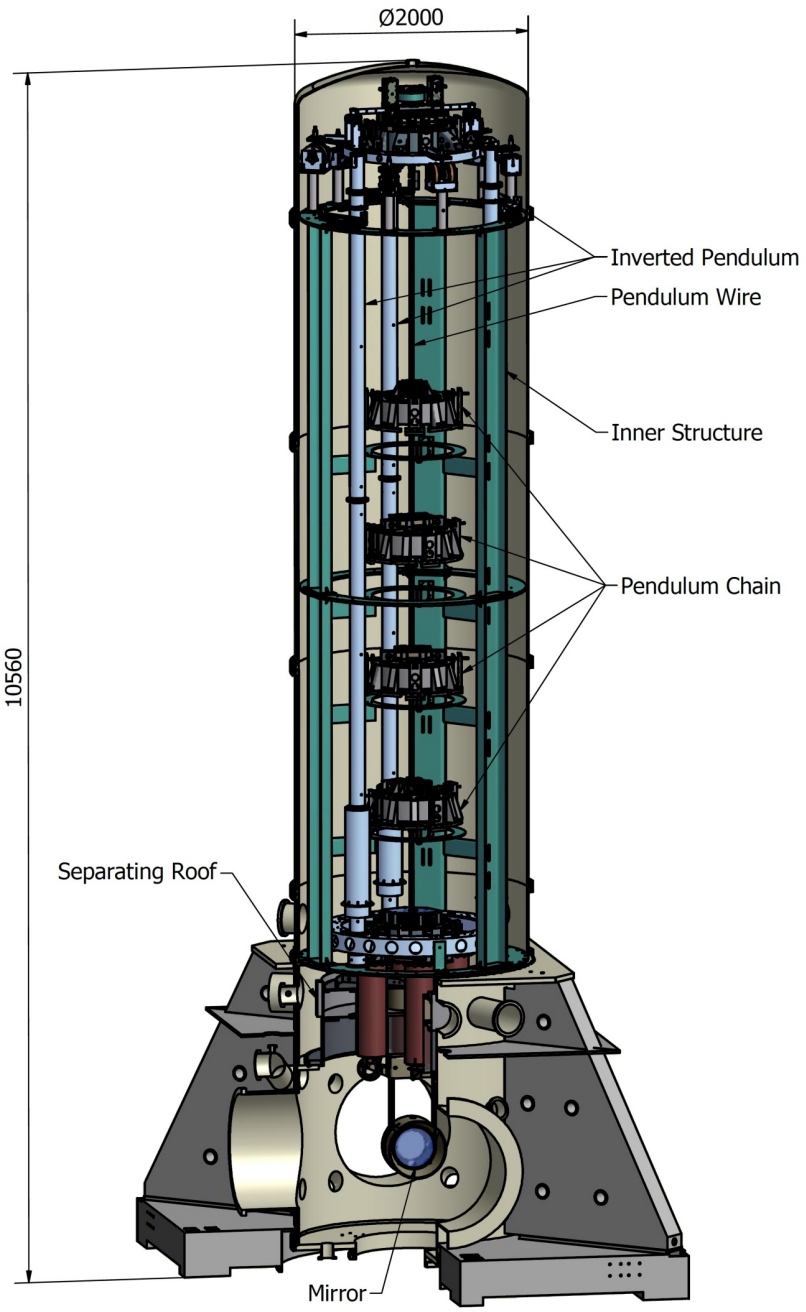

Figure 36. The VIRGO Superattenuator. The three-legged metallic structure of the Inverted Pendulum, the six seismic filters and the payload of a long suspension chain are shown inside their long tower.

mechanical couplings between different degrees of freedom. Thus, vertical ground motion will cause a change of the laser beam optical path.

Following all these considerations, it is clear that a vertical attenuation of seismic noise comparable with the horizontal one is fundamental.

With a multistage pendulum this goal can be achieved by replacing each suspension wire with

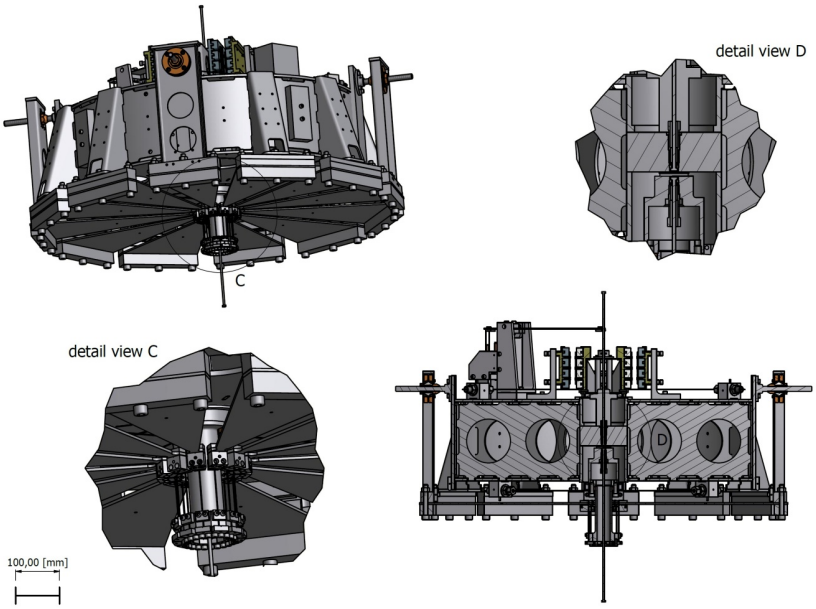

Figure 37. Technical drawing of a mechanical seismic filter.

a spring to form a cascade of vertical oscillators. The spring should support a heavy load and, at the same time, it should be soft enough to exhibit a low resonant frequency. In this way it is possible to confine the vertical resonances of the chain within a low-frequency region obtaining a strong attenuation, starting from a few Hz, as occurs for the horizontal pendulum resonances. Helicoidally-shaped springs may have internal low frequency resonances that would spoil the SA isolation performance. This problem has been overcome by using normal metallic suspension wires, attached to triangular cantilever springs clamped below each pendulum mass (figure 37).

With this technical solution the masses of the SA act as 'mechanical filters' in the vertical direction.

Also, the unavoidable mechanical couplings between rotation of the pendulum masses and horizontal mirror displacements can cause optical path changes. To confine these rotational mode frequencies below the detection band, each pendulum mass has been designed to have high momentum of inertia. In addition, the diameter of the suspension wire, connecting two consecutive stages, has to be small enough to reduce its restoring torque, determining the rotational frequencies. Furthermore, upper and lower wires have to be attached as close as possible to the centre of mass of the filters, assuring low frequency rotational modes and negligible coupling effects on horizontal displacements of the test mass.

A drawback of such a complex system is that seismic motion is amplified in correspondence to its resonance modes. This is cured by active damping, as described in the section on SA control.

6.3 The Superattenuator design

The Virgo SA, which represents one of the peculiarities of the apparatus, was designed by taking advantage of the working principle of a multistage pendulum [26]. In addition to the pendulum

2012 JINST 7 P03012

chain, each SA includes a 'top stage' and a 'payload'. The top stage consists of a three-legged inverted pendulum with a top platform. The payload consists of an upper mass, called a 'marionette', supporting an optical component. The optical component may be either a mirror (a test mass) or an optical bench with several optical and electronic elements. In the mirror case, the marionette also supports a so-called 'reference mass', suspended behind the mirror and acting as a reaction mass for the actuators controlling the mirror position.

Two differently sized SA have been designed to fulfil the different needs involved in isolating mirrors and optical benches. Each mirror is suspended to a long SA chain with 6 filters (figure 36), while a bench is suspended to a short chain with only two filters (figure 38).

Long and short SA are accommodated inside towers 11 m and 7 m high, respectively.

Here below we describe a long SA, while additional information on each single chain element can be found in reference [27].

6.3.1 The Inverted Pendulum top stage

The top stage of the SA is designed to fulfil three main functions:

- to introduce an additional very low frequency ($\sim$30 mHz) horizontal filtering stage;
- to provide the SA with a suspension point positioning system: the tidal strain over 3 km, of the order of one hundred microns over a period of 6 hours, is far beyond the dynamic range of the actuators exerting the steering forces upon the mirror. Therefore tidal drifts are compensated by gently moving the SA suspension point at the top of the seismic attenuation chain;
- to provide a soft suspension stage on the top of the chain allowing active mode damping and seismic noise reduction by means of inertial sensors, position sensors and electromagnetic actuators. In addition, any noise introduced by acting on the top stage is filtered out by the SA chain.

To fulfil all of the above requirements in the horizontal plane a device based on the principle of the inverted pendulum (IP) was created. An IP is a suitable device for several reasons:

- an ideal inverted pendulum can be conceived as a mass-less vertical bar of length l connected to ground by means of an elastic joint with stiffness k and supporting a mass M on its top. In such a pendulum the gravity acts as an anti-spring and the resonant frequency, given by

$$f_{IP} = \frac{1}{2\pi}\sqrt{\left(\frac{k}{M} - \frac{g}{l}\right)}$$

  can be tuned to be very low (about 30 mHz in Virgo);
- the force required to displace the chain suspended from an IP resonating at 30 mHz is very low: the force required to move the 1 tonne SA chain by 1 cm is less than 1 N. Therefore soft electromagnetic actuators can be used to control the mirror position. For this reason the IP is a good platform to act upon for the active damping of the SA normal modes.

2012 JINST 7 P03012

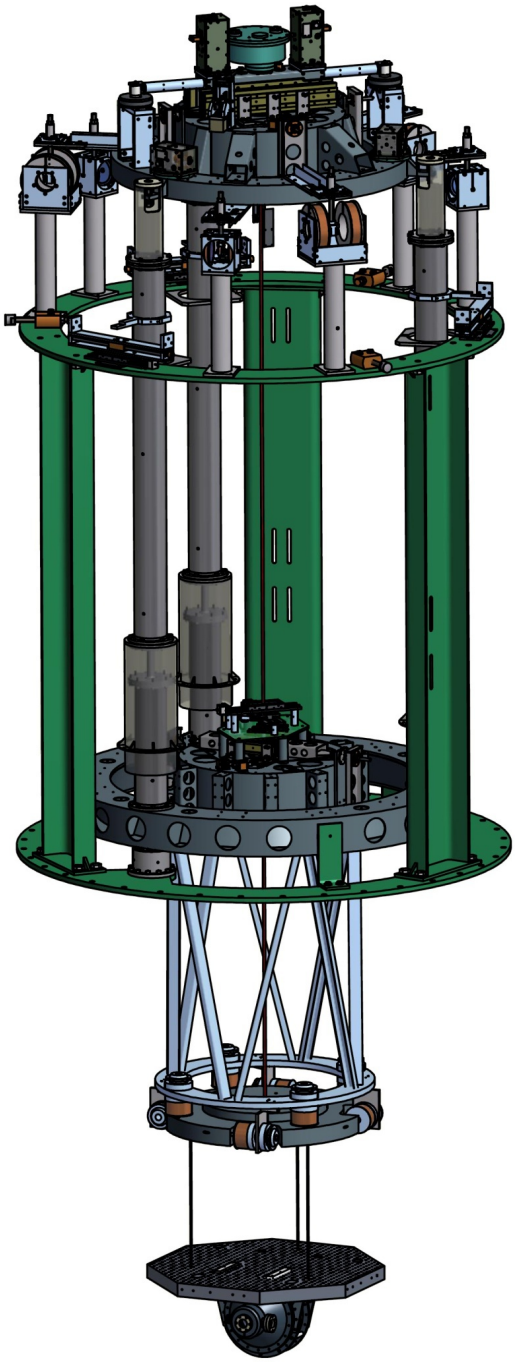

Figure 38. One short Superattenuator. This figure is relative to the Injection Suspension where a marionette with the shape of a disk is shown, as well as the optical bench under which the reference cavity for the laser beam is attached.

The Virgo IP (figure 36) consists of three metallic legs clamped through flexible joints to a bottom ring resting inside the base of its vacuum tower. A steel ring is suspended to the top ends of the three legs, by three thin wires (31 mm long). This ring surrounds the first filter of the chain (hereafter called Filter Zero) to which it is rigidly connected. The Filter Zero and the top ring

2012 JINST 7 P03012

together form a platform to which the whole pendulum chain is suspended. The bottom ring is not bolted to the tower but rests on three hollow steel feet, to be used for the installation of piezoelectric actuators below, in order to correct for seismic ground tilts.

Each leg is a 26 kg aluminium pipe with an inner diameter of 125 mm and a 2.5 mm thickness. The total length of the leg, from the bottom of the flexible joint to the suspension point of the SA, is about 6.2 m. The leg is composed of three sections, flanged together and reinforced with titanium inserts.

Due to the non-zero mass of the leg, an extension below the flex joint is necessary to tune the percussion point [28], to preserve the isolation performance. A proper counterweight is attached to a bell-shaped 0.9 m-long skirt bolted to the bottom of the leg. The flex joint is therefore mounted on top of a 0.8 m-high rigid support.

The IP structure is surrounded by a rigid metallic structure (called the inner structure in figure 36), holding the parts of sensors and actuators that need to stay 'on ground'. It is equipped with [28, 29]:

- 3 motorised sleds on the external structure in front of the leg tops. Each sled is connected to the corresponding leg top through a soft spring acting in a direction tangent to the top ring. The motors are used to set the IP roughly in its working point;
- 3 horizontal LVDT (Linear Variable Differential Transformer) position sensors set in pin-wheel configuration. The secondary windings stay on the inner structure, while the primary ones are rigidly connected to the top stage;
- 3 horizontal coil-magnet actuators, set in pin-wheel configuration. The coils, arranged as Maxwell pairs, stay on the inner structure, while the magnets are rigidly connected to the top stage.

6.3.2 The seismic filters

As mentioned above, in the SA each pendulum mass is a vertical mechanical filter. A sequence of six filters is able to isolate the optical components from the vertical seismic noise with a working principle similar to that of a multistage pendulum. A detailed description of the design and performance of a single filter can be found in reference [30].

A seismic filter is a rigid stainless steel cylinder (70 cm in diameter and 18.5 cm high, for a total weight of about 100 kg) suspended as close as possible to its centre of mass (figure 37). A set of triangular cantilever spring blades is clamped on the outer circumference of its bottom part. Each blade (3.5 mm-thick and 385.5 mm-long) is bent at a constant curvature radius and with different base width according to the load to be supported. The different curved blades become flat and horizontal under loads ranging between 48 and 96 kg. Each blade tip is connected by a 1 mm diameter steel piano wire to a mobile central column, coaxial to the filter body. Any movement of the central column, apart from the vertical direction, is prevented by two sets of four centring wires stretched on the top and on the bottom of the filter body.

A crossbar, bolted on the upper part of the central column, is used as a mechanical support for the magnetic anti-spring system described in the next section. The central column and the crossbar represent the moving part of the mechanical filter from which the load of the lower stages

2012 JINST 7 P03012

is suspended by a steel suspension wire. By connecting each filter to the next one, a chain of mechanical oscillators in the vertical direction is obtained.

The set of triangular steel blades mounted under each filter, sustains the weight of the lower part of the chain (figure 37). The support capability of each spring set is tuned varying the number of blades from 12 to 4, the base width of each blade between 180 mm and 110 mm and the curvature radius without load.

Once properly loaded, the main vertical resonance frequency of each filter is about 1.5 Hz. In order to reduce the peak amplitude of the blade first flexural mode (about 100 Hz), a rubber/metal damper is mounted in the middle of each blade. Another special damper, mounted on the crossbar, suppresses a mode at about 50 Hz, due to the suspension wire connecting the crossbar to the next filter and acting as a stiff spring [30].

'Maraging©' steel has been used in place of standard steel for blade construction minimising micro-creep effects [31, 32] due to the high load applied. The same material was also used to machine the suspension wires with nail-heads at both ends [33]. Since the suspended load decreases, moving from the top to the bottom of the chain, the wire diameter changes along the chain in the range 4-1.85 mm. In this way it has been possible to confine the wire violin modes within the high frequency band and to reduce their angular stiffness, which determines the rotational frequencies around the vertical axis.

The two nail-heads of the wires connecting a filter to that above and below are screwed into the central part of its body at a relative distance of 5 mm, very close to the filter centre of mass. As mentioned above, this guarantees a small return torque to rotations of the filter around a horizontal axis and thus a low tilt frequency.

6.3.3 The magnetic anti-spring

The chosen suspension wire length of 1.15 m sets the pendulum resonant frequency of each stage at about 0.5 Hz. In the vertical direction, the stiffness of the triangular blade springs fixes the natural resonant frequency at about 1.5 Hz.

In order to reduce the vertical stiffness of the blades, and then to confine the main vertical resonant frequency of each filter below that of the pendulum, a system of 'magnetic anti-springs' [30, 34] was conceived. This consists of four sets of permanent magnets (two assembled on the crossbar and two on the filter body) facing each other with opposite horizontal magnetic moment (namely in a repulsive configuration). The two matrices screwed on the crossbar move in vertical direction only.

The working principle of this system is sketched in figure 39.

When the magnets are perfectly faced, the repulsive force has a null vertical component, but as soon as the crossbar (and its magnets) moves in the vertical direction, a vertical component of the magnetic force appears. Considering a small relative displacement (Δy), the vertical component of the repulsive force (F_y) is proportional to Δy:

$$F_y \approx F_0 * (\Delta y / d)$$

where F_0 is the module of the repulsive force and d is the magnet separation (figure 39). Such a device is equivalent to a vertical spring with a negative elastic constant (anti-spring) whose module

2012 JINST 7 P03012

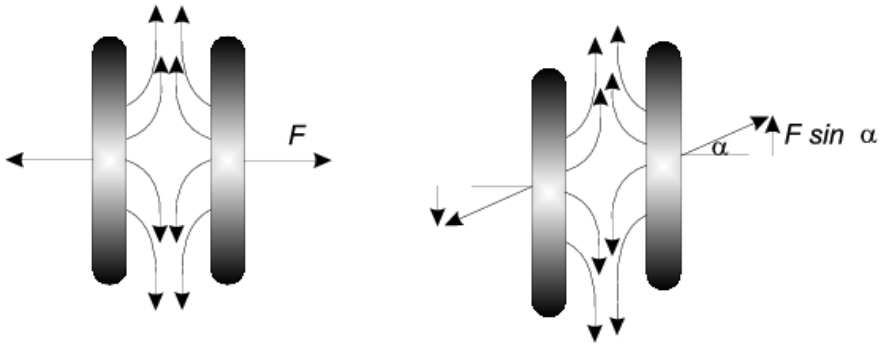

Figure 39. The working principle of the magnetic anti-spring system. When the magnets are displaced in the vertical direction (right-side) a vertical component of the repulsive force appears.

is F_0/d. Its rest position is thus the position for which the two couples of matrices are perfectly faced.

On a seismic filter the magnetic anti-springs act on the crossbar in parallel with the blade mechanical springs, so that the overall elastic constant is reduced and the vertical mode frequencies of the chain are confined below the highest frequency of the horizontal modes.

The device working principle requires that in the rest position all the magnetic anti-springs are perfectly aligned, in order to null the vertical component of the magnetic force. To this end a small remotely adjustable cantilever blade (the so called 'fishing rod') has been added on top of each filter, to bring the anti-springs to the correct working point.

6.3.4 The Filter 7

The last mechanical filter of the chain is the so called Filter 7 [35]. Since it has been designed to steer the payload, it has been equipped with four aluminium legs (about 900 mm-long and 250 mm in diameter) bolted under the filter body bottom part. These legs are used as a mechanical support for the coils mounted in front of the permanent magnets screwed on the marionette wings. The coil-magnet pairs are used to apply forces and torques to the marionette, controlling its position in three degrees of freedom: translation along the beam direction (z), rotation around the vertical (y) and transverse axes (x).

In addition, the Filter 7 is equipped with different vacuum-compatible stepping motors. Two of them are dedicated to the filter alignment around its vertical axis. The first one is accommodated on top of the filter body and is used to rotate the filter with respect to the upper part of the SA, also moving the whole payload. The second is mounted under the filter and rotates the marionette and the rest of the payload with respect to Filter 7, adjusting the relative positions of coils and magnets.

A second pair of vacuum-compatible stepping motors moves two small masses attached to mechanical trolleys on top of the filter body, performing a fine adjustment of the tilts (rotations around the two horizontal axes).

The set of sensors and actuators is completed with three positioning sensors with permanent magnets screwed on to the cylindrical filter body mantle. These are installed reproducing the geometry of the IP top stage (pin-wheel configuration at 120°C from one other) and, together with

a set of three 110 mm diameter Maxwell pair coils, mounted on the ground connected bottom ring, represent the system used to feedback control the Filter 7 position during interferometer alignment. The large size of the coils, ensuring a field gradient uniformity of better than 2% over a 30 mm distance ensures that the noise reintroduced at this stage is negligible.

6.4 The mirror payload

The mirror payload, comprising the marionette, the mirror and its reference mass, is suspended under Filter 7. The mirror and its reference mass are independently suspended to the marionette. The payload was designed to provide two separate levels of action to control the final mirror position: Filter 7 on the marionette and reference mass on the mirror. The actuators consist of coil-magnet pairs. The mechanical resonance frequencies of the payload have to be as high as possible to limit thermal noise and to reduce cross-talk between the various degrees of freedom to be controlled.

The marionette is composed of a body and four arms carrying the small permanent magnets of the electromagnetic actuators. The marionette body is made out of stainless steel AISI 304L and its mass is 108 kg. The four arms, are made out of non-magnetic steel AISI 316 L, to avoid spurious magnetisation effects.

The marionette is suspended from the SA by means of a 1 m long Maraging©steel wire with 1.8 mm diameter. This wire is clamped to the marionette centre of mass in order to minimise the torques needed to control the angular position. The vertical position of the attachment point can be adjusted with an accuracy of the order of 0.1 mm. A motorised mass of about 240 g housed in the lower part of the marionette allows to adjust the inclination around the $\boldsymbol{x}$-axis within ± 7 mrad. In this way the marionette can compensate the mirror θ_x offset.

The mirror rests on a pair of thin steel wire loops, descending from the marionette and transferring directly to the mirror the steering action made by Filter 7 coils. The reference mass is supported in an identical way and has two notches to allow the passage of the mirror support wires, without any contact.

The reference mass has the shape of a hollow cylinder, coaxial with the mirror, which is contained within it. In this way, the two centres of mass are coincident. This enables a reduction of the coupling between the various degrees of freedom when a force is applied on the mirror by the reference mass actuators. Moreover, the reference mass has the function of protecting the mirror and its delicate coatings from contamination, mainly during assembly, and to avoid any mirror damage by supporting it in case of wire breakage; The present reference mass is made out of Al alloy 5056, for an easier machining. It is planned to replace it with a dielectric reference mass to avoid eddy currents.

All the payload components are ultra-high vacuum-compatible and the reference mass shape limits the dust contamination of the mirror during the assembly phase and in operation.

The marionette is steered by means of the coil-magnet actuators along four degrees of freedom: translation along the vertical (y) and optical (z) axes, rotation (θ_x) around the transverse horizontal axis and (θ_y) around the vertical axis (figures 40, 41, 42). Another set of coil-magnet actuators permits the application of forces directly between the reference mass and the mirror, along $\boldsymbol{z}$, $\boldsymbol{x}$, $\boldsymbol{\theta_x}$ and $\boldsymbol{\theta_y}$, for automatic alignment and locking purposes. To this end, the reference mass is equipped with four coils facing four small magnets glued to the rear face of the mirror and two coils facing the magnets on the mirror sides.

2012 JINST 7 P03012

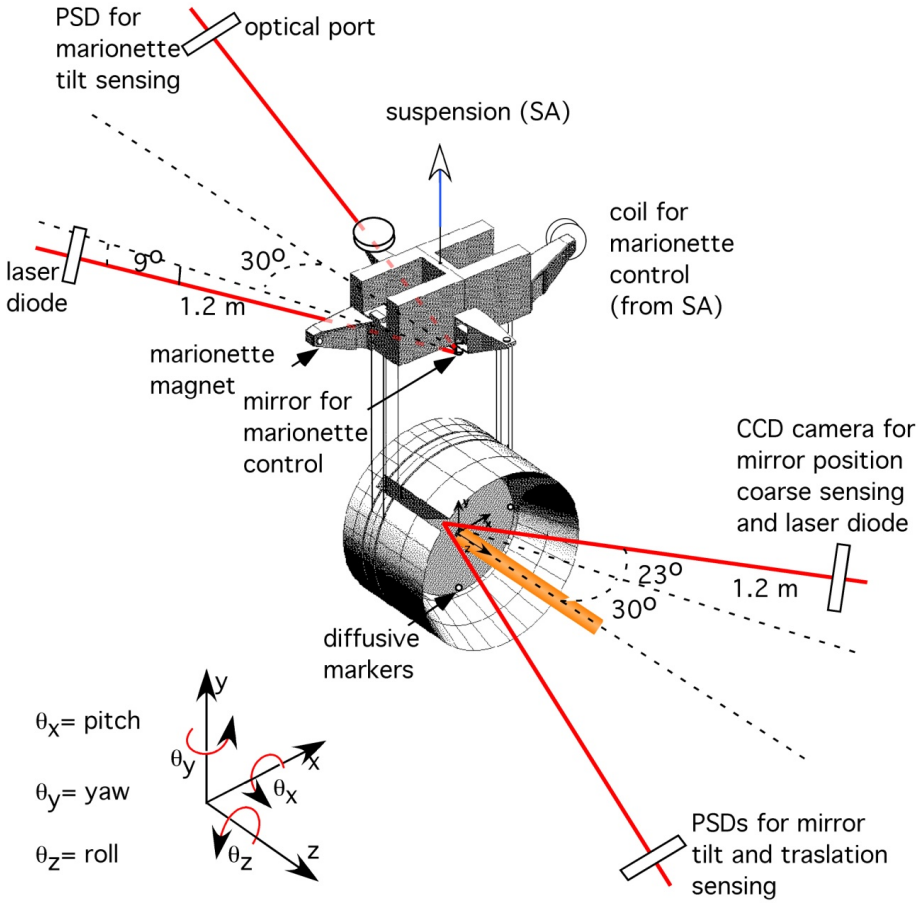

Figure 40. The mirror payload and the standard reference frame used to describe the motion of the optical elements of the Virgo interferometer.

6.4.1 Mirror suspension wires

Present limitations in interferometer sensitivity below 100 Hz are mostly due to thermal noise from the mirror supporting wires [36–38]. The noise spectrum is given by the Fluctuation-Dissipation theorem [39] that relates the position fluctuation spectrum $x^2(f)$ to the system dissipation or, in the following expression, to the pendulum transfer function $H_{pend}(f)$:

$$\left(x^2(f)\right) = \frac{4k_B T}{2\pi f} \Im\left\{H_{\text{pend}}(f)\right\}$$

where k_B is the Boltzmann constant and T the system temperature. The imaginary part of the pendulum transfer function contains the dissipation terms that determine the quality factor Q. A full treatment of the thermal noise behaviour of the Virgo suspension can be found in [40].

For the mirror wires C85 steel was selected as the best compromise [41] between internal dissipation and breaking strength requirements.

Relaxation of wires under tension ('creep') could be a source of mechanical shot noise. This was investigated in [42] and a thermal treatment has been developed to reduce the noise contribution: the suspension wire is left under tension at 150°C for one week in a 'U' shape and in a Nitrogen atmosphere, to facilitate further mirror assembly.

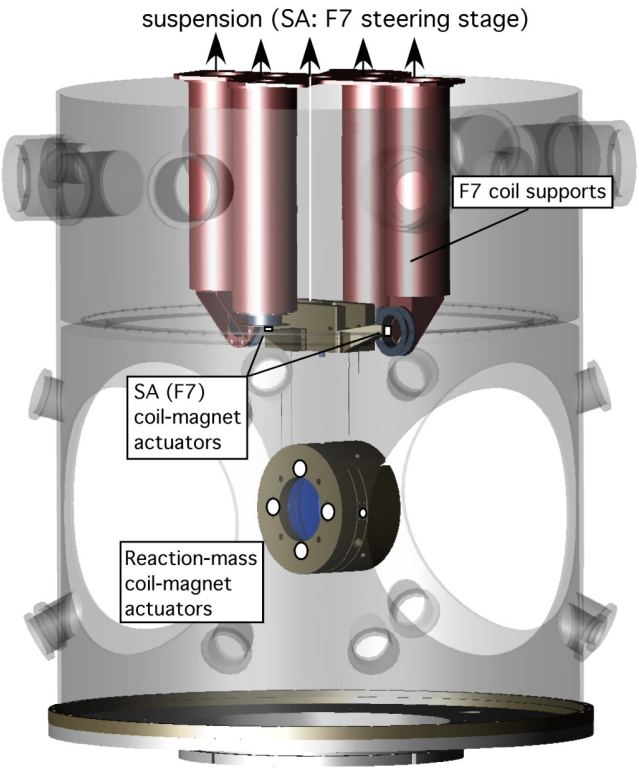

Figure 41. 3-D representation of positions of coils and magnets for the local controls.

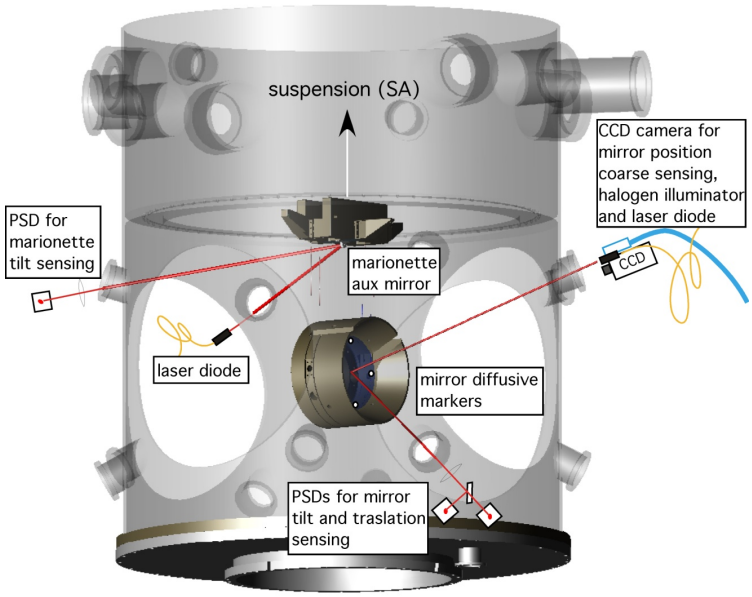

Figure 42. 3D representation of sensors for the local controls.

2012 JINST 7 P03012

In addition, dissipation occurring in the wire clamping system of the marionette and at the contact point on the side of the mirror must be taken into account. For this reason the clamps in the marionette have been designed [43] in such a way as to minimise stick-and-slip processes, which could increase mechanical energy dissipation: the wire is clamped into a stainless steel clamp through two hardened steel inserts, in which a groove has been incised through an electro-erosion process. The clamping pressure is then controlled, in the assembling phase, through a dynamometric wrench. To reduce friction between the suspension wires and the lateral surface of the mirror, a fused silica prism (a so-called 'spacer') is bonded to the lateral surface of the mirror. Silicate bonding [40, 44] is used to preserve the mirror mechanical quality factor. The performance of this suspension has been evaluated in [36] and a quality factor of about 10^6 is expected.

6.4.2 Payload assembly

Payload assembly is a complex procedure performed in a class 100 clean room and respecting the strict requirements relating to accuracy, in terms of mirror and reference mass positioning with respect to the marionette, and to mirror handling and cleanliness.

To assist in the reaching of this aim, a stainless steel assembly structure was built. This structure was equipped with 4 micrometric sleds, so as to be able to adjust the position of the reference mass and of the mirror with a mechanical accuracy of 0.1 mm along the translation degrees of freedom x, y, and z. Similar adjusting gears are used to control the tilt of the marionette, reference mass and mirror.

Before the assembly of the whole system both the reference mass and the marionette were equipped with coils and motors, and electrically cabled.

The fused silica supports for magnets were attached to the mirror rear face by silicate bonding. The magnets were then fixed with a UHV-compatible ceramic glue. On the front face of the mirror, near to its edge, where the reflective coating is absent, four markers for mirror position measurement were bonded.

At this point the real assembly work started, requiring about one day per mirror. The most delicate phase was when the mirror was taken out of its class 1 box and placed on a cradle below the marionette. Then, a sucking disk flange was positioned in front of its coated side using the positioning gear. Contact on the uncoated external area was achieved through a circular gasket. After evacuation, the mirror supports were released. In this way the lateral surface of the mirror was made available to position the suspension wires.

The reference mass was then put on to a cradle and positioned around the mirror. Its suspension wires were then clamped to the marionette.

At this point, all of the elements of the payload were secured to the assembly structure and protected from dust with a multi-layer hermetic plastic cover.

The payload was then lowered through a trap-door from the clean room floor into the gallery, which leads to under the suspension towers. An air cushion trolley was then used to take the payload to the proper tower. The tower trap-door was then opened and clean air flushing started. The trolley arrived under the tower and the payload was raised. Then two operators entered the tower to suspend the payload and make the electric connections. The end mirrors were taken to the end buildings through the 3km tunnel using a low pressure tyre electric trolley.

The suspension procedure, involving two operators in the tower, lasted about 4 hours. The most delicate part of the operation was when the payload was balanced and rotated in order to be aligned with the coils attached to Filter 7 and the upper suspension was tuned.

6.5 The bench payload

The payload of the short suspension was equipped with a marionette with a different geometry. In this case, the Filter 7 legs were replaced by an aluminium structure (see figure 38) supporting eight coils (four vertical and four horizontal). The coils are faced toward eight permanent magnets, which are screwed into a disk-shaped marionette (800 mm in diameter and 105 mm thick). With this set-up the coil-magnet system is used to control the payload in four degrees of freedom: the displacement along the beam axis (z), the rotation around this horizontal axis (θ_z), the rotation around the second horizontal axis (θ_x) and the rotation around the vertical axis (θ_y).

These short suspensions, only to be used for the Injection, Detection and Mode Cleaner payloads, were conceived with an Inverted Pendulum supporting a filter chain, with two mechanical passive filters. In the first two SA, the payloads consist of a marionette, with the above-mentioned geometry, and an octagonal optical bench, while in the Mode Cleaner SA, a small reference mass surrounding a curved mirror (80 mm in diameter) is hung to the marionette.

The different marionette geometry mounted on the bench payload was chosen as the short towers are not equipped with a separating roof dividing the vacuum volume into two sections: the upper part with a low vacuum level and a lower part with high vacuum environment dedicated to the high optical quality mirrors. This marionette also serves a protective purpose in relation to optical elements and electronic devices mounted on the benches.

6.6 Mirror control

The overall task of this system is to drive the interferometer mirrors from an incoherent to a coherent state, in which the mirrors are in a fixed relative position at the level of picometers and nanoradians, according to the Global Control correction signals. Due to the large position range of the Virgo suspensions, hierarchical sensing and actuation of mirror position is required. Concerning the operation of hierarchical control during the standard operation of the interferometer, the task is achieved through the performance of large corrections at low frequency, upstream from the mechanical filters. In this way, introduced control noise is filtered out and does not reach the mirror. Intermediate amplitude corrections at intermediate frequency are performed at the marionette level and only fine corrections are applied at the mirror level (figure 43) [45].

The control system has to handle the following situations:

- re-alignment of the suspension;
- pre-alignment of the interferometer;
- local angular control and damping of payload normal modes;
- lock acquisition;
- automatic alignment;

2012 JINST 7 P03012

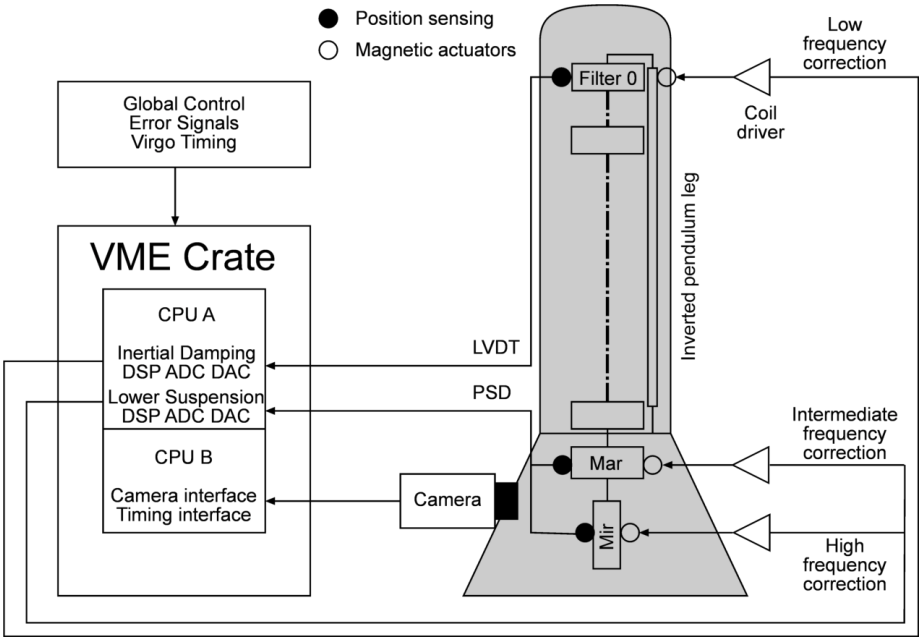

Figure 43. Hierarchical structure of the suspension Control System.

- locked standard operation.

The sensing and actuation devices used to control the Virgo mirror payload along 4 d.o.f. are integrated to constitute, together with the standard digital chain, the so-called Local Control (LC) system (figures 40, 41, 42) [46], a system which is replicated over 8 locations, i.e. 6 mirror payloads plus two suspended benches.

In the incoherent state the mirrors are controlled by sensors located close to the suspensions and referred to ground. Hence their relative positions are mechanically correlated by the ground and subject to independent drifts ($\sim 1\mu$ rad/h). In the coherent state the control is operated through the beam itself, picked-off in several locations, and the mirror positions are optically correlated by interferometer signals (Global Control). The system has large operational dynamics, as it enables the recovery of the mirror alignment from ±50 mrad off, and makes it possible to maintain it with an accuracy of < 0.2 μrad rms, as well as to damp mirror oscillations at the level of the ground seismic figure (10^{-7} μm/Hz$^{1/2}$ at 1 Hz). The system is designed to accomplish a highly automated level of operation. Software issues will be the subject of a coming paper. The control servo-loops, using local and ground-based mirror position read-out, have a bandwidth from DC to 2-3 Hz and are meant to be left open once the control is undertaken through signals derived from the coherent state of the interferometer (locked). The transition from local to global control of the interferometer has to be accomplished in a smooth way.

6.6.1 Mirror position measurement

Since LC sensors are clamped to the ground their accuracy has to be comparable with the level of the ground seismic noise; this is achieved using classical optical means. The overall sensing

2012 JINST 7 P03012

hardware, modified according to the experience acquired with the Central Interferometer [46], is described here-below (figures 40, 41, 42).

Coarse mirror position read-out is performed by a CCD-camera (EEV-17CAM), which is located in air at 1.2 m from the mirror and at 35° incidence angle. The CCD-chip is a scientific grade 512x512 pixel matrix on a 7x7 mm^2 area. The pixel clock is set at 20 MHz and the device provides a digital signal through an 8-bit ADC. The image is gathered and converted in 16 ms. A dedicated read-out board was designed to use this device in real-time applications in a standard VME environment [47, 48].

Fine mirror position read-out is provided by an optical lever, measuring the mirror position along three d.o.f.. This is composed by a 1 mW laser diode coupled through an SM-PM optical fibre (OZ Optics), to an anti-reflection (2 10^{-3} at 670 nm) coated viewport. Following reflection onto the suspended mirror, the beam exits the vacuum tank through a symmetrically positioned optical viewport, impinges on a 50% splitter cube and then hits two Position Sensing Devices (PSD Hamamatsu S2044, 4.7x4.7 mm^2) located at the focal and image planes of a lens with 200 mm focal length. The system is rather compact and the laser is injected through the same optical port used for the CCD camera. This two-sensor device provides a measurement in 3 d.o.f., i.e. pitch, yaw and longitudinal motion of the mirror along the optical axis.

Fine marionette control of angular pitch and yaw (θ_x and θ_y) is provided through a further optical lever, impinging at a 30° incidence angle on a 1" metal-coated auxiliary mirror attached to the bottom side of the marionette. The sensors used are of the same kind as those mentioned above.

Coarse and fine position signals are managed by a server running in a dedicated real-time CPU. Suitable error signals and logical flags are built and passed to the DSP (Digital Signal Processor) where digital filters are applied (figure 43) [46]. According to the hierarchy of the actual architecture, the marionette-level read-out is the last error signal servoed before passing the angular control to the automatic alignment based upon wave-front sensing.

The optical components of each local control system are located on 30x30 cm^2 removable optical benches clamped to the suspension vacuum chambers inside the tower baking ovens. Each optical element of local control read-out is equipped with fine adjustment stages. This enables the possibility to reach optical diagonalisation of the read-out signal (i.e. separating mirror translations from tilts and pitch from yaw) reducing cross-talk between d.o.f. to 0.1-0.5 %.

6.6.2 Marionette and Mirror Actuators.

Several requirements were taken into account during the design of the actuators that are used to control the Virgo payload (see section 6.4). Some of these are related to the geometrical constraints imposed by the marionette and reaction mass shapes and to the force needed during the procedures to control the mirror position. The other requirements are dictated by the need to keep the injected electromechanical noise below the intrinsic noises of the system and to keep the power dissipated into the coils as low as possible.

To keep the interferometer locked it is also necessary to act on the mirrors in the detection frequency band, above 10 Hz. This modifies the interferometer response and, therefore, the actuators, together with their electronics, need to be precisely calibrated. Due to the required accuracy, the calibration cannot be performed on each suspension separately. Instead, the interferometer it-

2012 JINST 7 P03012

Table 3. Actuators present in the last stage of each suspension. The content of the different columns is explained in the text.

n	Loc.	axis	d.o.f.	Mag [T]	Coil Turns Turns	R_{coil} [Ω]	α [N/A]
2	M	Y	θ_x, z	1	1092	12	0.025
2	M	Z	θ_y, z	1	1092	12	0.025
4	m	Z	z, θ_x, θ_y	1	256	5	0.010
2	m	X	X	1	40	0.5	0.0053

self, controlled as a whole, must be used to determine the mirror actuator transfer function (see section 10.3).

All the actuators acting at the level of the suspension are coil-magnet pairs. Four coils control the marionette from the last seismic filter of the chain (Filter 7). In between the coils and the marionette magnets, metal-coated Kapton membranes (0.1 mm thick) separate the seismic attenuator vacuum chamber from the lower part, where the marionette is suspended. Another six actuators are designed to operate at the mirror level. Four coils are attached to the rear side of the reaction mass and act on the corresponding mirror magnets for position control and damping along the beam axis. These allow fast damping of mirror motion modes in the pitch (θ_x), and yaw (θ_y) d.o.f. as well. The last two actuators are rather compact and act on the sides of the mirrors along the x coordinate, in order to damp transverse motion. A detailed description of the actuator design can be found in reference [49, 50].

Table 3 summarises the actuators present in the suspension last stage. Cylindrical Samarium-Cobalt UHV-compatible permanent magnets (ϕ =11mm, h=4 mm for the marionette and ϕ =5mm, h=3 mm for the mirror) were chosen. Copper wires with a cross-section of 0.75 mm^2 were used to provide the required actuation current without overheating the payload. The coil-magnet nominal relative position is such that the actuator force response is maximum (z_{opt}=R_m/2, where R_m is the mean coil radius). The geometry of the coils was designed to be compatible with the mechanical room available, in order to widen as much as possible the flatness of this curve around the operation point [49, 50].

The number of actuators in each set is shown in column 1. The target object of the actuator is either the marionette (M) or the mirror (m). In column 3 the magnetic moment orientation is provided, using the Virgo standard reference system, shown in figure 41. The d.o.f. controlled by the given actuator is shown in column 4. Column 5 provides the magnetic field at the pole surface; column 8 the correspondence between current and applied force.

6.6.3 Suspension control electronics

Each of the 9 Virgo Superattenuators has a dedicated suspension control unit made of boards for data conversion, CCD camera image acquisition, processing and communication, which are located in a VME crate (figure 43). Each suspension uses around 60 channels of analogue I/O, 4 digital optical links (DOL, section 8.1.1), a CCD camera and a Timing system interface. Supervision and camera read-out software run on two PowerPC-based CPU boards.

Two DSP boards, based on the Motorola DSP96002 floating-point processors, compute the control signals using digital filters. These are triggered at 10 kHz by the timing distribution system to read error signals from ADC and DOL and to write correction signals back to DAC and to DOL. The digital filters used typically have 100 poles and each require a sustained rate of about 10 million floating point operations per second (MFlops). The overall computing power for the 9 Superattenuators is in the range of 200 MFlops.

The ADC boards contain 8 differential channels, which use 16 bits (corresponding to 14.5 effective bits) and are operated at 10 kHz. These are equipped with analogue anti-aliasing 7^{th} order Cauer filters with a 3.7 kHz cut-off frequency. The maximum signal on the differential inputs is 20 Volts peak-to-peak.

The DAC boards contain 8 channels, which use 20 bits (17.5 effective bits) and are also operated at 10 kHz. The analogue anti-aliasing filters are the same as those of the ADC boards. Balanced outputs provide a 20 Volt peak-to-peak signal. These components, which were not commercially available when the required specifications (speed and dynamic range) were established, were designed by our engineers.

The coils used for mirror and marionette actuation are fed with a suitable current, generated by coil driver circuits. These convert the voltage control signal into a current flowing in the coil. The coil drivers have two electronic channels in order to allow the use of large control signals during lock acquisition and reduced ones with a low noise level when the lock is acquired. These electronic channels are called High Power (HP) and Low Noise (LN), respectively.

Within a Suspension Control Unit, in addition to the available communication channels between boards, VME and VSB, a private protocol was developed to increase data throughput between DSP and conversion boards.

7 Optics

7.1 Introduction

As outlined in section 2.1 a complex optical system is used to measure the path difference, induced by gravitational waves, between the two arms of the Virgo interferometer. The required precision in the measurement is achieved by having a pure TEM00 monochromatic light beam and using the best available optical elements to preserve the beam quality. The error signals and finally the gravitational wave signal are obtained through the analysis of the light beams at the several output ports of the interferometer (figure 44).

This section first introduces the full Virgo optical scheme. The light source and injection system are then described, followed by the Virgo mirrors. Finally, the detection system, used to measure the output beams and to extract the relevant signals, is described to complete the section.

7.1.1 Virgo optical scheme

The laser beam used in the interferometer has to be stable in phase, lateral jitter and power [51]. This is achieved first through the use of a suspended Input Mode Cleaner (IMC) cavity, which performs many functions. The IMC length is used as a reference in the laser frequency prestabilisation circuit. The IMC also filters the input light, thus determining the position and angle of the beam that goes into the interferometer, keeping only the TEM00 component. The resulting intensity fluctuations are corrected acting on the laser power by means of a feedback loop.

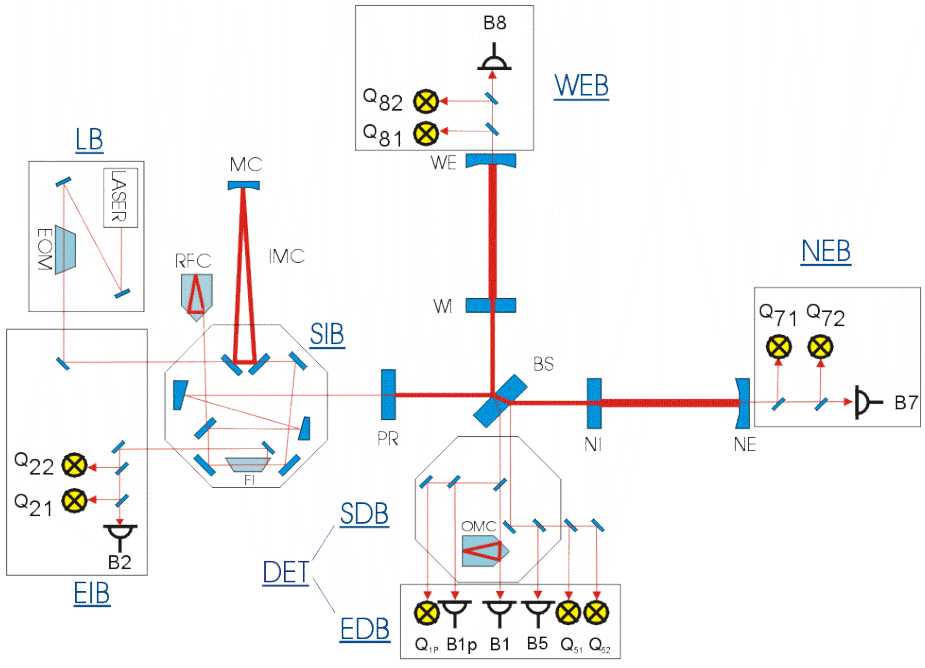

Figure 44. Virgo optical lay-out. Suspended Injection Bench (SIB), Input Mode Cleaner (MC) end mirror, Reference Cavity (RFC), North/West Input/End (NI, NE, WI and WE) mirror, Power Recycling (PR) mirror. From B1 to B8: interferometer output beams, analysed by the corresponding photodiodes. Electro-Optic Modulator (EOM).

2012 JINST 7 P03012

The beam entering the interferometer is divided by a 50% reflecting beam splitter, into two orthogonal beams. The light power in the arms is enhanced by Fabry-Perot cavities with a finesse of 50. A power recycling mirror sends the light returning towards the laser back to the interferometer, forming a power recycling cavity with almost 500 W incident on the beam splitter; the power in the arm cavities becomes 8 kW. The two arm beams recombine on the Beam Splitter. At the output port, looking at the second face of the BS, the interference dark fringe is analysed by a complex detection optical bench.

The five main beams exiting the interferometer (figure 44) are analysed with photodiodes, and provide information about the condition of the interferometer. As described later in Chapter 8, this information is used by the locking system to maintain the resonance in the arm cavities and in the recycling cavity, and to keep the destructive interference at the interferometer output. The correction signals are sent to the actuators controlling the positions of the end mirrors, the beam splitter and the power recycling mirror. The combined error signal of the arm cavities is used for a second level of laser frequency stabilisation. Moreover, each output port contains a set of quadrant photodiodes, providing information on the misalignment of the mirrors, which is used in the auto-alignment system to correct mirror angles and the input beam alignment. The Mode Cleaner and reference cavity have their own system of locking and auto-alignment.

7.2 The injection system

Due to the very stringent requirements in relation to the light source, the injection system has a rather complex lay-out, shown in figure 45.

7.2.1 The 20 W laser

The power laser (slave laser) is a 20 W, solid state laser: two Nd-YVO4 rods are pumped by two 30W fibred diodes. The cavity is a four mirror ring cavity in 'bow tie' configuration, the cavity support being made of invar, with 2 mirrors glued on to piezoelectric actuators. The optical/optical efficiency is about 50%.

The slave laser is injection-locked to a 'master' laser (1 W Non-Planar Ring Oscillator laser from Innolight Cie). As a result, the slave is phase-locked with the master laser, and therefore, the slave laser frequency can be driven by the master laser frequency [52]. The slave laser pumping diodes are inserted into a feedback loop to control the power being sent to the interferometer.

The injection locking range is about 3 MHz. Both lasers must be kept within this range for the injection process to be effective: this is obtained through piezoelectric actuators driving the slave laser cavity length, the achieved unity gain frequency being about 80 kHz. The associated error signal is based on the Pound-Drever-Hall technique: the master laser beam is phase modulated at 14 MHz and the error signal is extracted from the beat of the sidebands with the carrier at the reflection port of the slave laser.

This set-up enables the slave laser to work with full power (about 22 W) on a single longitudinal mode, the frequency fluctuations being those of the master laser itself ($10^4/f$ Hz/Hz$^{1/2}$), plus a residue, kept negligible thanks to the slave laser analogue loop and injection process. At the end of 2006 the system had been running continuously for more than 27000 hours.

Two Faraday isolators, one before and one after the Electro Optical Modulator (EOM), provide the attenuation needed by the laser system (more than 40 dB). The beam exiting the slave laser

2012 JINST 7 P03012

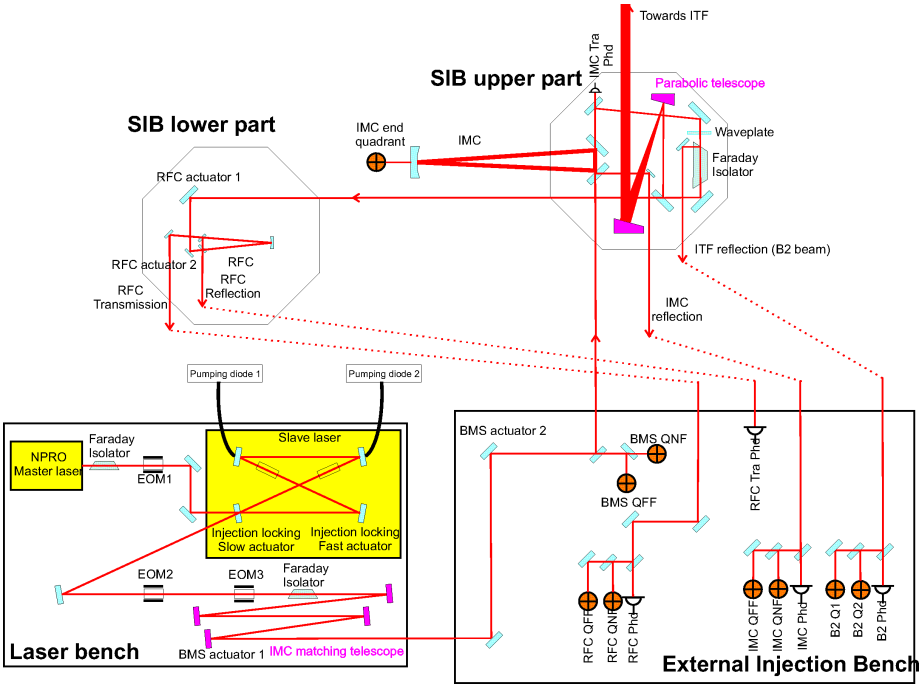

Figure 45. The Virgo Injection System. The Master and Slave Laser bench is on the bottom left. The light goes to the External Injection Bench on the bottom right. The Suspended Injection Bench top and down views are shown on the top right and left respectively.

is passed through a first Faraday isolator (LB_F2 in the figure), then collimated inside another EOM (LB_EOM2), which provides the phase modulations at 22 MHz, 6 MHz and 8 MHz. The first modulation is used for the longitudinal lock of the Input Mode Cleaner (IMC), while those remaining are used for the lock of the whole interferometer. The second Faraday isolator (LB_F3) is after the EOM. Downstream, an all-reflective telescope (LB_M9 – LB_M12) expands the beam from 2 mm diameter to 1 cm, as required to match the IMC cavity.

The whole laser system is accommodated on a dedicated optical bench (Laser Bench, LB), resting on ground within a plexiglass protection and an acoustic isolation enclosure.

7.2.2 External Injection Bench

The beam coming from the Laser Bench passes through an acoustically isolated tube to the External Injection Bench (EIB). The EIB rests on ground in the laser laboratory and is acoustically isolated. It supports the optics to steer the beam to the Suspended Injection Bench. In addition, it hosts the photodetectors measuring the light reflected by the injection system and the interferometer itself (beam B2).

Before entering the vacuum chamber of the Suspended Injection Bench, a small fraction of the input beam is sent to a Beam Monitoring System (BMS), based on two DC quadrant photodiodes and related optics. This system aims to correct beam drifts with respect to a ground-based reference (DC read-out photodiodes). One quadrant photodiode (BMS_QN) is placed on the image plane of the last mirror of the Laser Bench output telescope (LB_M12) and the other (BMS_QF) on the focal plane of a telescope, so that beam lateral shifts and angles can be controlled. The signals of

2012 JINST 7 P03012

these photodiodes are linearly combined in order to enable the performance of pure angular and translational motions with respect to the IMC input mirror. The actuation is obtained using two mirrors (LB_M11_ABP1 and EIB_M3_ABP2) controlled by piezoelectric crystals.

The beam reflected by the IMC is sampled by a broadband photodiode (IMC_DR) and is used for IMC longitudinal locking. The same beam is sampled by two broadband quadrant photodiodes (IMC_QF, IMC_QN) performing the wavefront sensing for the IMC automatic alignment, measuring position and angle respectively. In addition, a digital CCD camera images the beam.

In a similar way, the signals to monitor and control the Reference Cavity (RFC) are generated using a photodiode (RFC_DR), two quadrants (RFC_QF, RFC_QN) and a camera.

Finally, the reflected beam from the interferometer (B2) is sampled for longitudinal locking purposes using three photodiodes (B2_f_6MHz, B2_3f_6MHz and B2_f_8MHz). The signal from these photodiodes is demodulated using the 6 MHz reference signal, its third harmonic and the 8 MHz reference respectively. The beam position is measured using broadband quadrant photodiodes (Q21 and Q22) placed in order to have the required Gouy phase on the photodiode planes. Finally, a digital CCD camera images the interferometer reflected beam.

7.2.3 Suspended Injection Bench and Input Mode Cleaner

The Suspended Injection Bench (SIB) is an octagonal, 900 mm diagonal optical bench suspended in vacuum from a short Superattenuator. Its main components are the input optics of the IMC, the RFC, which is placed below the bench, a large aperture Faraday isolator, and the optics to match the beam to the interferometer and the related photodetectors.

Coming from the External Injection Bench, the beam enters the SIB tower through a large anti-reflection coated (AR) 3° inclined silica window, directly reaching the input mirror of the Input Mode Cleaner. The reflected beam is sent back to the External Injection Bench by means of a plane mirror (SIB_M8) placed on the SIB. The transmitted beam is sent through a Faraday isolator and an off-axis parabolic mirror telescope to the interferometer through a diaphragm located at the exit of the SIB. The telescope is necessary to widen the beam to the size required by the interferometer, equipped with a flat Power Recycling mirror. In the past, this was achieved taking advantage of a flat-concave Power Recycling mirror, with the drawback of being sensitive to transverse displacements. The Faraday isolator turned out to be necessary to prevent light coming back from the interferometer and entering the Input Mode Cleaner. These elements were not present in the first version of the SIB.

A small fraction of beam transmitted by the IMC is monitored by the photodiode (IMC_D1T) used by the laser power stabilisation system.

Downstream from the Faraday isolator a small fraction of the beam is sent to the Reference Cavity (RFC) by means of a periscope (SIB_M11-SIB_M12). The RFC is a 30 cm-long monolithic triangular cavity, attached beneath the injection bench. The body is made of ULE®[53], an ultra low thermal expansion coefficient ceramic, supporting three optically contacted mirrors. It is used as a frequency reference in the first stages of lock acquisition, so that the laser frequency line width is lower than the recycling cavity line width. Its finesse with P-polarisation is 1000, providing a DC frequency reference better than 10^{-2} Hz/$\sqrt{\text{Hz}}$ [54].

2012 JINST 7 P03012

The beams transmitted and reflected by the RFC are sent respectively to a photodiode placed on a small external bench at the level of the RFC, and to the RFC photodiodes and quadrants on the EIB. In this configuration the RFC cannot be locked before IMC locking.

The Input Mode Cleaner is an in-vacuum suspended triangular Fabry-Perot cavity with a length of 143.6 m. Its design finesse is 1000. The input mirrors (IMC1 and IMC2) are flat Suprasil substrates 80 mm in diameter and 30 mm thick, coated on one side for Nd:YAG at a 45° angle of incidence, S polarisation, and AR coated on the other side. The two mirrors are rigidly connected, being optically contacted to a pierced Zerodur [55] block forming a dihedron with an aperture angle of slightly less than 90°, in order to close the light path. The dihedron is resting on a three point support using steel stems mounted on the SIB. The cavity end mirror has a diameter of 80 mm, a radius of curvature of 180 m and is suspended in vacuum from a short SA. The incidence angle on the far end curved mirror is 300 μrad, reducing the astigmatism aberration losses. The beam waist is 5 mm and the light intensity 170 W/mm^2.

The IMC free spectral range is 1.044 MHz. This allows for a large choice of modulation frequencies for the main interferometer beam, of which they must be a multiple. The selected modulation frequency is 6.264 MHz. The achieved matching of the input beam to the IMC is of the order of 98%. However, owing to IMC internal losses, the transmission at the IMC resonance is of the order of 50%, thus yielding about 10 W maximum power.

Before entering into the interferometer the beam goes through a Faraday isolator (SIB_F) with a quite large aperture (20 mm). Even with this aperture, the beam coming out of the IMC has to be reduced to about half of its size (from 5 to 2.65 mm), in order to avoid being clipped.

Downstream from the Faraday isolator, an 8x magnifying optical system is used to match the beam size to the interferometer. This collimating parabolic telescope is formed by one 2" diameter +75 mm focal length mirror (SIB_M5), with a 12.7 mm off-axis, and a 4.5" diameter +600 mm focal length mirror (SIB_M6), with a 101 mm off-axis.

All of the optics are adequately dimensioned to also deal with parasitic beams, where required, and to avoid beam clipping. Behind some mirrors, PSD detectors (position sensing photodiodes, PSDi on the drawings) have been located in order to monitor possible beam drifts. Finally, several beam steering elements are remotely controlled. The most critical actuators operate in a closed-loop mode to allow controlled and reversible movements.

The alignment of the SIB is performed in two stages: the SIB is locally controlled using a CCD camera looking at a small plate supporting four markers and two mirrors placed in front of the dihedron. The coarse alignment of the SIB is performed using the markers, illuminated by a halogen lamp. The fine control is performed using two HeNe beams, coming from outside, and reflected by the mirrors into the camera, which detects their displacement.

7.2.4 Injection system control

The expected frequency stability given by shot-noise resolution is 10^{-6} Hz/$\sqrt{\text{Hz}}$. With an expected common mode rejection factor of better than 10^{-2} this frequency noise does not limit the interferometer sensitivity. However, as the IMC mirror positions determine the light beam parameters and phase, this stability can only be obtained by means of a feedback system acting on the suspended optical elements and on the laser itself (figure 46).

2012 JINST 7 P03012

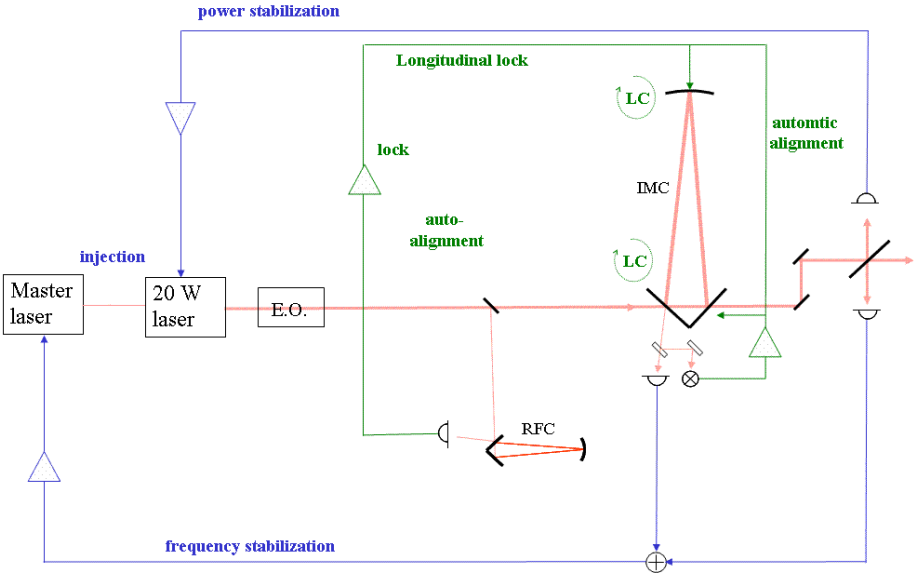

Figure 46. Injection system control scheme.

Due to the large dynamic range required for control, the laser frequency is pre-stabilised on the IMC. A Pound-Drever-Hall scheme with a phase modulation at 22 MHz is used, which analyses the beam reflected by the IMC. Before the full lock of the interferometer, the IMC length is stabilised on the RFC, with a 100 Hz bandwidth. The error signal is processed in a digital chain, and the correction signal is sent to the coils acting on the IMC end mirror.

The laser frequency is eventually stabilised using the interferometer common mode as a reference, with a relative frequency noise of 10^{-19} $Hz^{-1/2}$. This is called the Second Stage Frequency Stabilisation (SSFS). When the full frequency stabilisation is active, these two signals are added to provide the feedback error signal.

The laser phase is controlled by three different actuators: a corrective EOM device for fast actuation (30 kHz – 300 kHz), a piezoelectric crystal that controls the master laser cavity length at intermediate frequency, and the laser crystal temperature control (< 1 Hz).

The mutual alignment of the IMC and of the incoming laser beam is achieved with the Ward technique using two quadrant photodiodes, placed on the EIB. In addition, the beam position at the IMC end mirror is measured using a quadrant in transmission. These sensors provide six error signals, which include longitudinal and angular information. The corrections are sent to the SIB position control, as well as to the IMC mirror angular control. The IMC alignment correction is first engaged acting only on the end MC mirror, which therefore follows the beam. When the full alignment control of the IMC is switched on, corrections are also applied to the SIB and the Local Control can be switched off.

The power transmitted by the IMC is stabilised with a feedback on the slave laser pumping diodes. The feedback is AC coupled, with cut-off around 30 mHz and unity gain frequency at 40 kHz. The feedback is applied either on the DC current supply of the diodes or on a fast 'shunt' connected in parallel with one of the pumping diodes (fast actuation). In figure 47 the measured performance in terms of Relative Intensity Noise (RIN) is shown.

2012 JINST 7 P03012

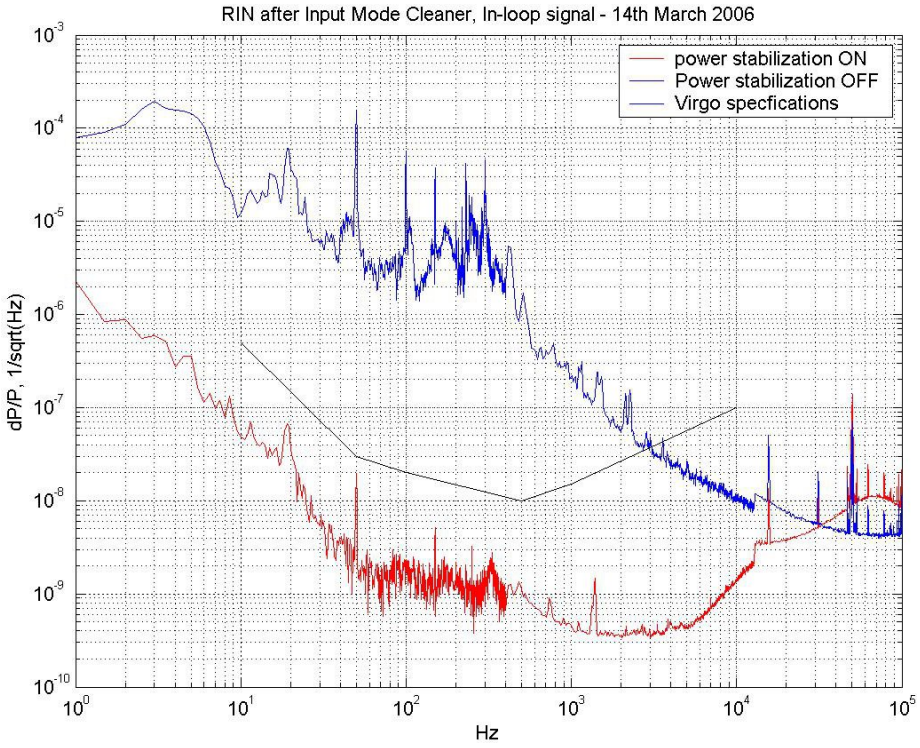

Figure 47. RIN of input mode cleaner transmitted beam, In-loop signal.

7.3 The detection system

7.3.1 The detection system design

The detection system analyses the beams coming from the output ports of the interferometer as shown in figure 48. The B1 beam provides the gravitational wave signal and is analysed by the main detection system. This system also measures the beam reflected by the second face of the beam splitter (B5), also providing error signals for the longitudinal and angular control loops. Beams B7 and B8 are analysed on benches located at the end of each interferometer arm, completing the set of control signals, and the beam reflected by the interferometer (B2) is analysed on the injection bench.

The main detection system is designed to measure variations of light power of the order of 10^{-11} W at the modulation frequency corresponding to changes in interferometer arm length smaller than 10^{-19} m. On the other hand, the mirror residual motion at low frequency, of the order of 10^{-12} m causes large fluctuations in the light intensity. This puts severe constraints on the dynamic range and residual noise of the photodiodes and their read-out system. Additional high quality optics are required to handle the various beams and to adapt their size to the photodetectors. The different components of the detection system are suspended in vacuum or located on an optical bench in air depending on their sensitivity to seismic and acoustic noise.

Due to various defects, such as mirror surface deformation and interferometer misalignments, higher order transversal modes are present in the B1 beam. This increases the contrast defect and can spoil the sensitivity by a factor of 2-3. An output mode cleaner in vacuum is used to improve this contrast defect.

2012 JINST 7 P03012

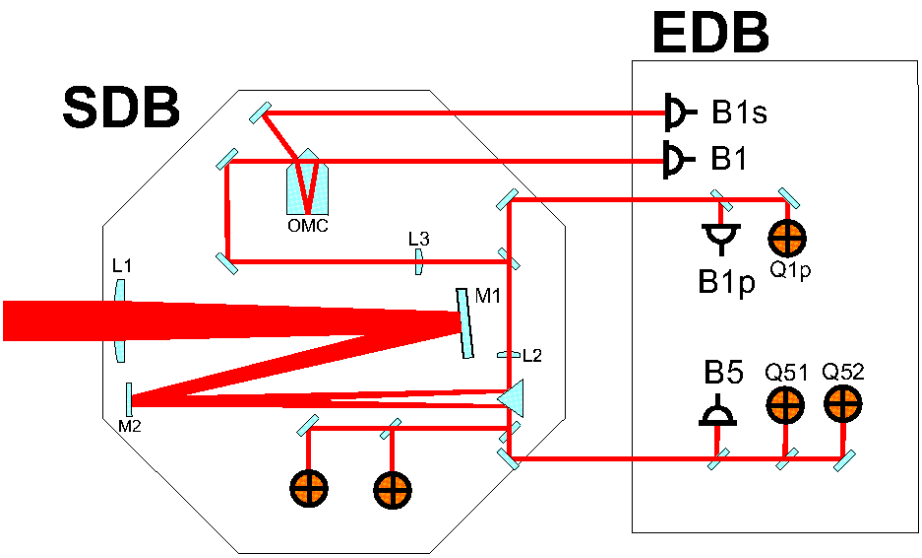

Figure 48. Schematic view of the detection system. B1 is the dark fringe beam and B5 is the beam reflected by the second face of the Beam Splitter mirror. B1p is a small fraction of B1 extracted before the output mode cleaner (OMC) and B1s is the beam reflected by the OMC.

The read-out photodiodes are located on a second bench placed outside the vacuum chamber and the produced signal is demodulated according to the Pound-Drever-Hall technique. A principle scheme of the main detection system is shown in figure 48. Other similar benches are also located at the transmission port of the Fabry-Perot cavities and at the reflection port.

7.3.2 The output mode cleaner

The output mode cleaner (denoted OMC on figure 48) is a short triangular monolithic cavity (2.5 cm-long) made out of silica with a finesse of 50 [56]. The large cavity bandwidth (75 MHz) allows for the transmission of both the carrier and the sidebands in the same Airy peak. Due to this choice a small change in the mode cleaner length induces a change in the relative phase of the carrier and the side-bands and simulates a gravitational wave signal. Therefore, its length has to be controlled with a precision of about $3\text{x}10^{-10}$ m. To reduce the effect of external disturbance, the cavity is suspended to a seismic isolator inside a vacuum chamber.

The cavity length is controlled by varying the OMC temperature using Peltier cells. A temperature scan makes it possible to identify the fundamental mode TEM_{00} by looking at the shape of the transmitted beam. The OMC is kept in resonance by modulating the cavity length at 28 kHz with a piezoelectric device and detecting the error signal synchronously. The achieved length precision is about 10^{-10} m (rms).

7.3.3 The suspended detection bench

The output mode cleaner is located on an optical bench suspended in vacuum to a seismic isolator. This bench also contains a telescope, the purpose of which is to adapt the interferometer output beam (B1) to the output mode cleaner cavity: the waist of the beam is reduced from 2 cm to 140 μm. The positions of the 3 lenses (L1, L2 and L3 on figure 48) forming this telescope can be

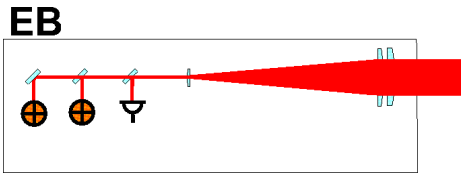

Figure 49. Schematic lay-out of the North and West End benches.

remotely adjusted in order to tune the waist of the beam and its location. The alignment of the beam on the mode cleaner is performed using the signals from two quadrant photodiodes (DQ1 and DQ2) and acting via pico-motors on the orientation of mirrors M1 and M2 located inside the telescope. Downstream from the output mode cleaner a Faraday isolator avoids that the small fraction of light back-scattered by the photodiodes recombines with the interferometer beam.

In order to keep the output mode cleaner aligned with the interferometer output beam, the suspended bench position must be controlled in all degrees of freedom with a precision of the order of 1 μm in translation and few tenths of μrad in angle. The control system [57] is based on a CCD camera placed outside the vacuum chamber. The camera looks at the image of four different clusters of LED, attached to the vacuum chamber and reflected by four mirrors attached to the bench. Using the image positions of the four clusters the bench position in all its six degrees of freedom is deduced. Each cluster is composed of 10x10 LED to increase the measurement sensitivity. The camera noise enables the measurement of displacements of a few 10^{-8} m/Hz$^{1/2}$ and misalignments of a few 10^{-7} rad/Hz$^{1/2}$. The bench position measured by the camera is used as an error signal for a feedback performed using a DSP board and DAC channels driving eight electro-magnetic actuators. Once the feedback is closed, all the mechanical resonances are damped and the bench is controlled with the required precision.

7.3.4 The external detection bench

The external detection bench (figure 48), located inside an acoustically isolated enclosure, supports the photodiodes used to detect the gravitational wave signal (B1) and to provide the longitudinal and alignment error signals. It is equipped with CCD cameras to image the beams and monitor their position. The typical beam detection lay-out includes two photodiodes and, depending on the beam, quadrant photodiodes, CCD cameras and photodiodes demodulated at twice the modulation frequency. To avoid burning the photodiodes in case of large intensity fluctuations, a fast shutter is positioned on the B1 path.

7.3.5 The end benches

Acoustically isolated benches are located in the end buildings in order to detect the B7 and B8 beams transmitted by the Fabry-Perot cavities. Similarly to the external detection bench, each bench supports the photodiodes, quadrant photodiodes and CCD cameras. The photodiodes provide a low sensitivity length error signal, while the quadrants are used to control the alignment. The size of the beam coming out of the end mirror is adapted to the bench optics using a large lens doublet (see figure 49).

2012 JINST 7 P03012

7.4 The photodiodes and read-out electronics

7.4.1 Photodiodes for the dark fringe and the longitudinal control

Once the dark fringe is filtered by the output mode cleaner, the light is detected by photodiodes. Since the shot noise limit decreases as the square root of the detection efficiency, very high quantum efficiency photodiodes have been selected [58, 59]: these are 3mm diameter InGaAs photodiodes that, thanks to anti-reflective coating, have an efficiency larger than 90%. The diameter is the maximum imposed by the junction capacitance, in order not to limit the electronic bandwidth. These photodiodes are also used for the detection of the other beams used to longitudinally control the interferometer. Two photodiodes are used for each beam for redundancy. In case of bad contrast defects, up to 16 photodiodes are available for the dark fringe.

In order to detect the large signals at low frequency and the tiny signals at high frequency, the electronics must have a dynamic range of at least 10^8 and a noise equivalent to few 10^{-11}W/ Hz$^{1/2}$. The photodiode RF current is pre-amplified (with a typical gain of 400 V/A) using a filter which selects the signal at the modulation frequency. The noise of the preamplifier is 4 nV/Hz$^{1/2}$. Then the RF signal is demodulated yielding two output signals, the in-phase and in-quadrature signals, which are digitised with 18 bit ADC sampled at 20 kHz. With these ADC it is possible to achieve a dynamic range of 10^7 Hz$^{-1/2}$ (10 V peak over a noise of 1μV/Hz$^{1/2}$). In order to fit the ADC range the signal low frequency component is attenuated by a low noise analogue filter before digitisation. This increases the dynamic range of the detection system up to 2x10^8. Digital processing recovers the original signal.

For longitudinal control or interferometer characterisation purposes, the signals at 2 or 3 times the modulation frequency are also measured with photodiodes and electronics chains similar to those described above. The wrong polarisation beams are also extracted using polarising prisms and measured with dedicated photodiodes.

7.4.2 Photodiodes for the alignment control

The photodetectors used for the alignment control [60] are silicon quadrant photodiodes (EG&G YAG-444) with 11.4 mm diameters. The circular sensitive area is divided by two 125 μm gaps into 4 quadrants of 25 mm^2 each. The vertical/horizontal alignment error signals are obtained by building the intensity difference on the two halves of the detector (inside the front-end electronics box), followed by an off-board demodulation. The error signals are shot noise limited for an incident light power of 0.3 mW or more, the maximum allowed light power being 3 mW. The phase and quadrature signals contain information about angle and position of the wavefront of the carrier with respect to the modulation sidebands, from which the mirror misalignment can be deduced by using a suitably computed reconstruction matrix. This matrix can also include the additional information from the non-demodulated difference signals.

7.5 7.5 Mirrors

7.5.1 Introduction

The Virgo mirrors must meet very stringent specifications in order to allow the interferometer to reach the required sensitivity. The Virgo interferometer design requires very large (350 mm diameter) and heavy (20 kg) optics. The total losses of substrate and coatings at 1064 nm (including

Table 4. Requirements for the substrates of the Virgo mirrors.

Mirror substrate specifications				
Component	Fabry-Perot End Mirrors	Fabry-Perot Input Mirrors	Power Recycling Mirror (as of 2006)	Beam Splitter Mirror
Material	Herasil 1 Top	Suprasil 312 SV	Suprasil 312 SV	Suprasil 311 SV
Diameter (mm)	350 - 0, +0.5	350 – 0, +0.5	350 – 0, +0.5	230 – 0, +0.5
Thickness (mm)	96 – 0.5, +0	96 – 0, +0.5	96 – 0, +0.5	55 – 0, +0.5
Shape Side B	Concave R = -3450+/-100 m	Flat	Flat	Flat
Shape Side A	Flat	Flat	Flat	Flat
Wedge (mrad)	1.1 +/-0.1	no wedge (<< 1µrad)	no wedge (<< 1µrad)	1.1 +/-0.075
Roughness Side A	1-1.5 Å rms	< 0.5 Å rms	< 0.5 Å rms	< 0.5 Å rms
Roughness Side B	< 0.5 Å rms	< 0.5 Å rms	< 0.5 Å rms	< 0.5 Å rms
Flatness Side B	8 nm rms $\varnothing = 150$ mm	8 nm rms $\varnothing = 60$ mm	8 nm rms $\varnothing = 60$ mm	13 nm rms $\varnothing = 120$ mm
Flatness Side A	50 nm rms $\varnothing = 200$ mm	8 nm rms $\varnothing = 60$ mm	8 nm rms $\varnothing = 60$ mm	13 nm rms $\varnothing = 100$ mm

absorption, scattering and large-scale wavefront distortion) should not exceed 100 parts per million (ppm). There are specific constraints on the absorption (< 5 ppm) due to thermal lensing, that is mirror heating and deformation due to power absorption, which causes wavefront distortion. The large angle scattering level has to be less than 5 ppm, to minimise noise due to scattered light, which might recombine with the output beam.

These specifications can only be met with very stringent requirements on substrate composition and high quality polishing. The detailed requirements on substrates and coatings are presented in tables 4 and 5 respectively.

A new type of silica was specially developed by the German company HERAEUS in collaboration with Virgo.

For the coatings it was essential to build new equipment in LMA. New tools were also developed to manipulate the mirrors between every step of the production process (cleaning, coating, metrology) in clean conditions with a total safety. These tools are also used on the Virgo site in Cascina.

2012 JINST 7 P03012

Table 5. Requirements for the coatings of the Virgo mirrors.

Mirror coating specifications				
Component	Fabry-Perot End Mirrors	Fabry-Perot Input Mirrors	Power Recycling Mirror (as of 2006)	Beam Splitter Mirror
Incident angle	0°	0°	0°	45°
Coating diameter (mm)	330	200	200	210
Coating Side B (S polarisation)	H.R. R>0.99995	AR $R < 5.10^{-4}$	H.R. $R = 95\%$	$T = R = 0.5$ +/- 0.005
Coating Side A (S polarisation)		$R = 0.88$	AR $R < 5.10^{-4}$	AR R<10^{-3}
Losses (absorption = scatttering)	< 5 ppm $\varnothing = 150$ mm	< 5 ppm $\varnothing = 60$ mm	< 5 ppm $\varnothing = 60$ mm	< 5 ppm $\varnothing = 120$ mm
Flatness Side B	8 nm rms $\varnothing = 150$ mm	8 nm rms $\varnothing = 60$ mm	8 nm rms $\varnothing = 60$ mm	13 nm rms $\varnothing = 120$ mm
Flatness Side A	50 nm rms $\varnothing = 200$ mm	8 nm rms $\varnothing = 60$ mm	8 nm rms $\varnothing = 60$ mm	13 nm rms $\varnothing = 100$ mm

7.5.2 Mirror substrates

In order to have a very low bulk absorption level (< 1ppm/cm), a very high refractive index uniformity and a low birefringence, the substrates were made out of a new type of fused silica (Suprasil 311 SV, Suprasil 312 SV). The OH content is very low (<<50 ppm), the refractive index is homogeneous in all directions and the birefringence is very low (< 5.10^{-4} rad/cm). The result obtained was satisfactory with an achieved absorption = 0.7 ppm/cm for both Suprasil 311 SV and 312 SV.

An American company (General Optics) polished the bulk parts. This company was the only one able to polish very large substrates, while at the same time having a low micro-roughness (< 0.5 Å rms), a good flatness (< 8 nm rms on Ø150 mm) and few punctual defects. All the polished components, that were delivered, complied with the Virgo specifications.

The design coating (quarter-wave stacks) of each mirror depends on its function. For example, the end mirrors have a very high reflectivity level, while the input mirrors have only 88% of reflectance; in order to obtain a finesse of the Fabry-Perot cavities close to 50.

7.5.3 Mirror coatings

For the coating process it was decided to use SiO_2 for the low index layer and Ta_2O_5 for the high index layer, as these have the best optical properties in the mid-infrared region. The extinction coefficient at 1064 nm is:

$$\text{for } SiO_2 : k = 4.010^{-7} +/- 0.5.10^{-7}$$
$$\text{for } Ta_2O_5 : k = 4.510^{-7} +/- 0.5.10^{-7}$$

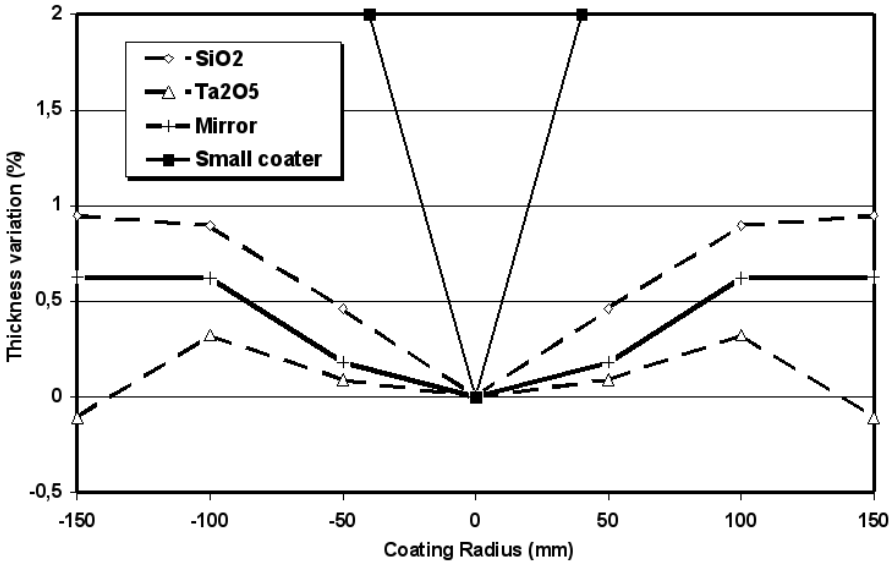

Figure 50. Thickness variation, with respect to the centre, as a function of the coating radius.

The quarter-wave multilayer mirrors resulting from this deposition system have an absorption of 0.6 ppm at 1064 nm. This value is lower than the Virgo requirement (5 ppm).

The SiO_2 layer is the restrictive factor regarding the layer thickness uniformity. Indeed, the uniformity is worse by a factor of 3 than that of the Ta_2O_5 layer (the emission of the sputtered particles is different for the two materials). The single layer thickness uniformity is shown in figure 50. Nevertheless, the uniformity obtained on a multilayer mirror was good enough for Virgo: 0.6 % on ∅300 mm.

The only deposition technique able to produce components with the required optical performance is Ion Beam Sputtering (IBS). LMA-Lyon have already demonstrated that this technique allows to reach very low-loss levels (< 5 ppm) [61, 62]. These were obtained on high reflectivity mirrors up to a diameter of 100 mm.

To reach the severe Virgo requirements, up to a diameter of 350 mm, a new Ion Beam Sputtering (IBS) coater was designed and manufactured. In order to obtain the lowest coating absorption, it was necessary to optimise the deposition conditions (beam current, target angles, gas flows). This was done by working first on single layers.

At the same time, LMA developed metrology tools that were able to characterise the optical performance of the coatings over a large area: wavefront, scattering, absorption and roughness measurements, as well as point defect detection. All these devices were installed in a new building in a class 1 clean room (figure 51).

This environment is necessary to avoid any contamination (particles) of the substrate before or after coating, which may increase the scattering and the absorption level.

7.5.3.1 Ion Beam Sputtering (IBS) coater

The new IBS system design was defined and optimised with developed-for-purpose simulation software [63], which precisely describes the sputtering phenomenon in the IBS process. This

2012 JINST 7 P03012

Figure 51. General view of the class 1 clean room (surface: 150 m^2).

software takes into account all of the geometrical and electrical parameters.

The mirror dimensions necessitated a vacuum chamber of 2.4*2.4*2.2 m^3 (figure 52):

This is the largest IBS coater in Europe. The sample holder (robot), operating under vacuum, allows to position the substrate anywhere inside a circle of 1 m. The coater was installed in a 150 m^2 class 1 clean room in 1999.

Primary pumping reduces the pressure to 8.10^{-3} mbar in 30 minutes, while the secondary pumping, performed by four CTI cryopumps, allows to reach a base pressure of 4.10^{-8} mbar in 3 hours. This pumping configuration guarantees a very clean vacuum, without hydrocarbon contamination, controlled in real time with a residual gas analyser (RGA).

Radio frequency ion sources are used to sputter the coating materials out of ultrapure targets (Ta_2O_5, SiO_2). Several reasons justified this choice instead of classical Kaufman ion sources. On the one hand, the possible layer contamination due to the tungsten filament, which degrades the layer absorption, is avoided. On the other, since the deposition time is very long (40-50 hours) and the filament life-time is comparable, the choice of filament-less sources was imperative.

The layer thickness is monitored by four quartz balances which allow for a very efficient in-situ control. Another component developed by LMA is the automatic control of the coater. This manages the pumping and venting procedure (done with ultrapure dry air) and every step of the deposition process (control of the ion sources, the gas flows, the target movements).

7.5.3.2 The optical metrology devices

The optical metrology of the coatings is fully part of the mirror production process. Indeed, it makes it possible to rapidly improve the coating performance and to reach a very high quality level. To avoid contamination of the mirrors, all the metrology tools are installed in the class 1 clean room, close to the IBS coater.

2012 JINST 7 P03012

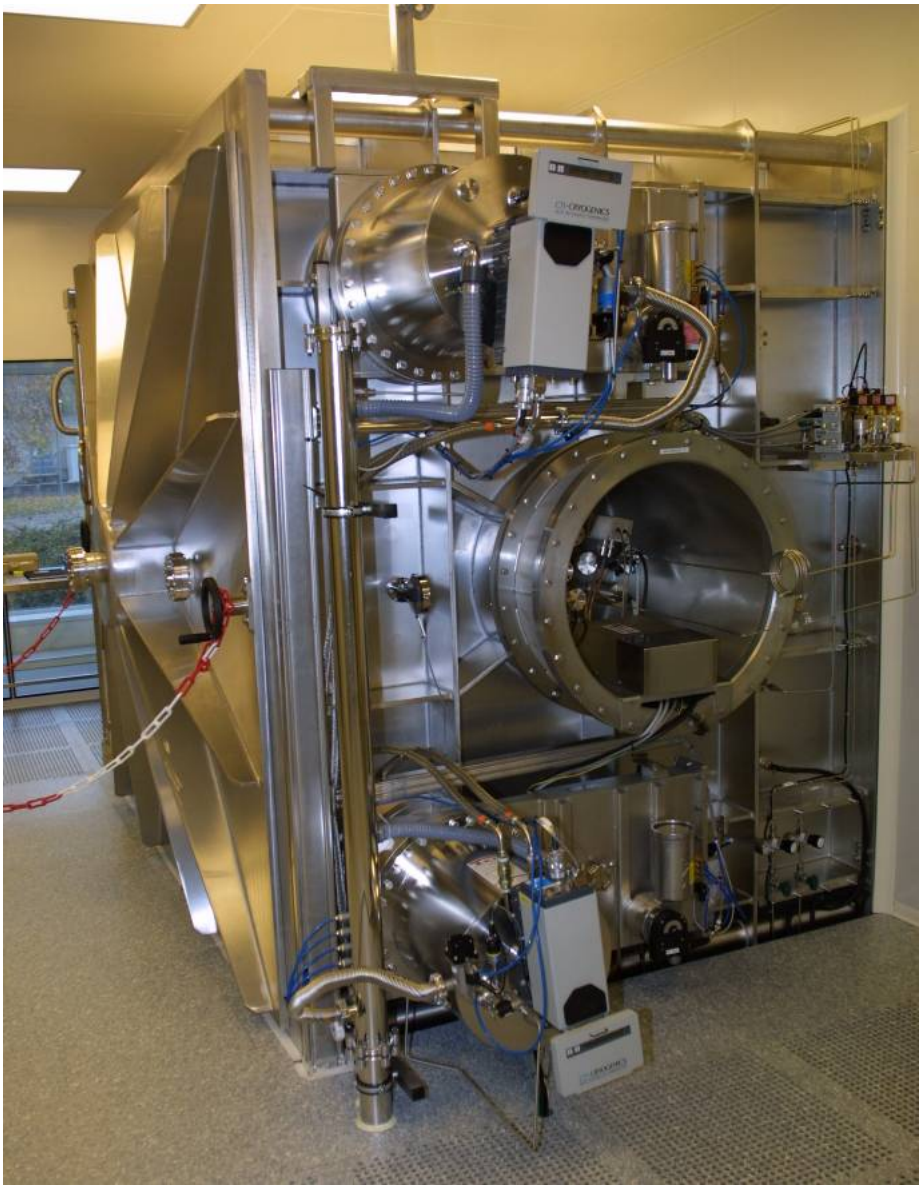

Figure 52. IBS deposition system.

All the metrology benches are able to characterise the entire surface of the large Virgo components (400x400 mm^2) and to handle heavy samples (20 kg). In this paragraph, we report only the main performances of the optical benches. More details can be found in [66].

Scattering and Transmission measurements: the scatterometer is a commercial CASI-type instrument able to measure scattering levels lower than 1 ppm (Bi-directional Reflectance Distribution Function noise floor of 10^{-9} sr^{-1}). The absolute BRDF as a function of the scattering angle θ_s is measured with a beam size of 1 mm over the whole sample. This system can also be used as a very sensitive wattmeter to determine very low transmission or reflectivity values (T < 1 ppm).

2012 JINST 7 P03012

Absorption measurements: a new absorption bench was developed to measure surface and bulk absorption on large samples. It is based on the photothermal deflection technique, the capability of which to determine very low absorption levels has already been demonstrated [61]. The sensitivity of this bench at 1064 nm is remarkable: 20 parts per billion (ppb) for surface absorption and 30 ppb/cm for bulk absorption have been achieved, making it possible to verify the high quality of the mirrors.

Wavefront measurements: a commercial interferometer from ADE Phase Shift (minifiz-type interferometer, pupil ∅150 mm, wavelength 1064 nm) is used. To measure large wavefronts (∅400 mm), a numerical method called 'Stitching Interferometry' [64, 65] was developed in collaboration with a French company (MB Optique). The reproducibility of this rather complex method, based on the difference between two successive measurements, is 0.4 nm rms on ∅400 mm. Without stitching the reproducibility is 0.04 nm rms on ∅150 mm.

Roughness measurement and defect detection: a commercial profilometer (MICROMAP), modified to characterise large optics, is used. The sensitivity is 0.2 Å rms. The roughness of all the Virgo substrates, polished by the American company General Optics, was found to comply to the requirement (< 0.5 Å rms).

The scattering level depends not only on the Rayleigh diffraction, which is linked to the surface roughness, but also on the number and the size of the point defects on the surface. To detect small defects (detection threshold ∅0.3 µm) a prototype based on the profilometer was developed. It is possible to check the substrate surface quality (maps on ∅400 mm) and also the substrate cleaning procedure, which is a crucial part of the deposition process. Small defects (between 0.3 and 1 µm), larger defects (> 1µm) and scratches are distinguished.

7.5.4 The Virgo large mirrors

The Virgo interferometer uses six large mirrors, as shown in figure 44. The substrate and coating specifications are listed in tables 4 and 5, while the measured optical performances are listed in table 6. The first set of Virgo mirrors was ready in June 2002. Figure 53 shows a Virgo mirror in its characterisation mounting.

The absorption and scattering requirements (< 5 ppm) are satisfied for each component. The dispersion in the absorption values comes from the coating design: the higher the reflectivity, the lower the absorption. Some scattering values are higher than 5 ppm. However, inside most of the central area, where the beam is more intense, scattering was measured to be within the Virgo specifications.

Another remarkable result concerns the wavefront flatness of these large components. The '8 nm rms on ∅150 mm' requirement on the wavefront is always satisfied. The factor limiting the reaching of the specifications is the substrate. Indeed, the rms flatness value is the same with or without coating (3.37 nm rms and 25.2 nm PV, 3.38 nm rms and 25.5 nm PV, respectively). The coatings reproduce the substrate surface exactly.

2012 JINST 7 P03012

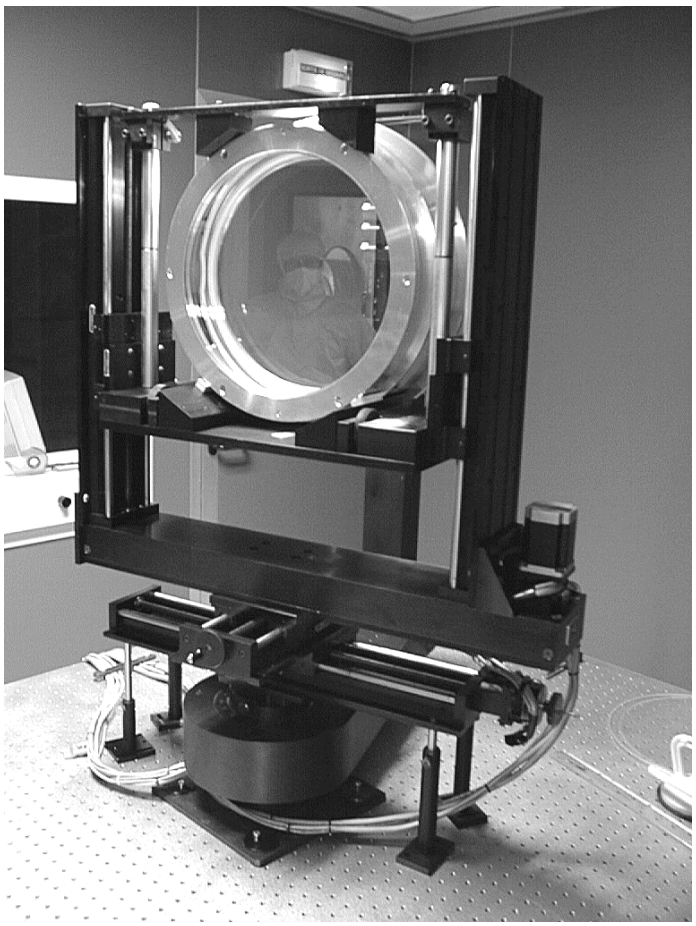

Figure 53. First Virgo End Mirror during the scattering measurement.

Table 6. Measured optical performances of the Virgo mirrors.

Measured optical performances of the large mirrors							
	Coating Nature and Clear aperture (mm)		rms Wavefront (nm)		Average Absorption (ppm)	Average Scattering (ppm)	Transmission
	Side A	Side B	Side A	Side B			
North End Mirror	-	HR ∅330 mm	3.9 ∅150 mm	3.8 ∅150 mm	0.67 ∅150 mm	4 ∅150 mm	42.9 ppm
West End Mirror	-	HR ∅330 mm	2.8 ∅150 mm	3.4 ∅150 mm	0.69 ∅150 mm	6.5 ∅150 mm	38.3 ppm
North Input Mirror	HR ∅200 mm	AR ∅200 mm	2.6 ∅60 mm	3.8 ∅60 mm	1.25 ∅60 mm	5 ∅60 mm	11.80 %
West Input Mirror	AR ∅200 mm	HR ∅200 mm	2.6 ∅60 mm	2.5 ∅60 mm	1.20 ∅60 mm	8 ∅60 mm	11.66 %
Recycling Mirror 2006	HR ∅200 mm	AR ∅200 mm	1.4 ∅60 mm	1.08 ∅60 mm	0.54 ∅100 mm	8 ∅60 mm	5.13 %
Beam Splitter n°1	AR ∅210 mm	HR ∅210 mm	3.9 ∅100 mm	5.4 ∅120 mm	1.35 center	5.5 ∅100 mm	49.8 %

2012 JINST 7 P03012

8 Overall interferometer control

8.1 Overall control

An apparatus like Virgo requires several control loops to reach and stay at its working point. Initially, the various systems (mainly the injection system, the mirror suspensions and the detection system) are run and controlled in an autonomous way (see dedicated sections for details). In order to reach the interferometer working point, the state of the various components must be steered in a global way. In its nominal running, two main feedback loops are active. The longitudinal control of the mirror positions (z degree of freedom as shown in figure 41) is achieved through the locking feedback loop. The alignment loop controls the transverse degrees of freedom (θ_x and θ_y in figure 41).

Figure 54 presents a schematic view of the different systems which are involved in the overall control of the interferometer:

- the timing system (section 8.1.2), which distributes a time reference over the site;
- the DAQ system, which collects all interferometer measurements. This is involved in the control loops as data provider for the Automatic Locking Procedure (ALP), section 8.3;
- photodiode and quadrant read-out (section 7.4, figures 55, 56), which collects the photodiode and quadrant signals;
- suspension control (section 6.6), in charge of the actuation on the mirror;
- laser and injection control (section 7.2.4);
- Global Control (section 8.2), which computes locking and alignment corrections;
- ALP, which provides a set of high level commands to connect the different systems (see section 8.3 for details).

The control system is implemented using digital servo loops. The locking loop runs at 10 kHz and the alignment loop at 500 Hz. Overall synchronisation is ensured by the timing system described in section 8.1.2 and the communication between systems is performed using Digital Optical Links (DOL) [70] (section 8.1.1). Locking and alignment loops behave in a similar way. The photodiode signals (both DC and AC, figure 44) are acquired by the photodiode and the quadrant read-outs (for locking and alignment, respectively) which send them to the Global Control. Using these signals, the Global Control computes the four lengths and the fourteen angles to be controlled. After appropriate filtering the Global Control sends the correction signal to be applied to each mirror suspension. At this level, depending on the state of the lock acquisition procedure, the corrections can be mixed with those coming from the local control system. This complex procedure is driven by the ALP which has access to the data collected by the DAQ system, allowing the determination of the interferometer status. When the science mode is reached, all mirrors are controlled in a global way and almost all the local control loops are switched off as they are noisier than the global ones.

2012 JINST 7 P03012

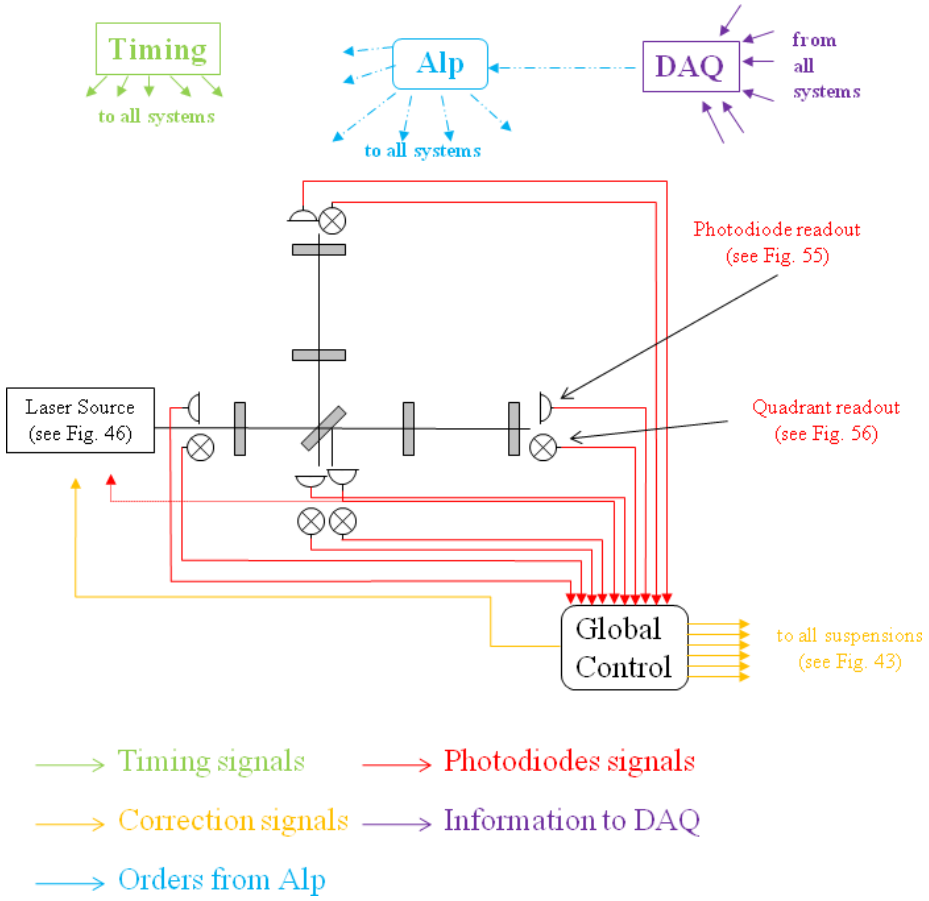

Figure 54. Overview of the systems involved in the overall control of the interferometer.

8.1.1 The hardware environment

Virgo decided to use the VME and VSB standard for the boards dedicated to the on-line system. The standard CPU for Virgo is the RIO8062 from CES [67], based on the PowerPC 604 chip from Motorola. This board was chosen for its good interrupt time response and its block transfer capabilities on VME and VSB [68]. The Virgo operating system is LynxOS [69], a UNIX-compatible real-time operating system for embedded applications. The version used is v2.5.1.

The various subsystems of Virgo (Photodiode Read-outs, Suspensions, Frame Builder, Global Control) only exchange digital data through optical fibres connected to a DOL-dedicated board. The DOL board is a VME/VSB slave module which is able to send a small number of data (some kBytes) in point-to-point mode, over long distances (up to 3.2 km), with no latency. Moreover, powering off of the receiver DOL does not affect the transmitter DOL, which enables a decoupling between the subsystems. It is designed with 2 channels: one to send data, one to receive data, which are completely independent. On each channel, a FIFO buffer takes care of the asynchronous operation between the optical link and the local bus of the board. The chosen FIFO buffer depth is 512 32-bit words for the transmitter channel and 512 or 16K words for the receiver channel.

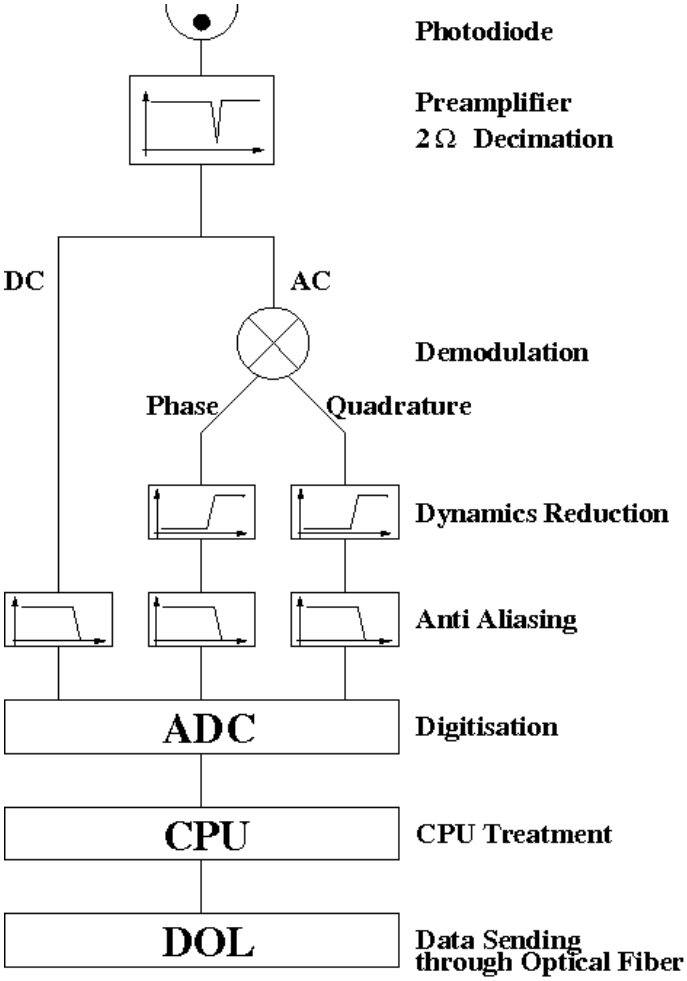

Figure 55. Photodiode read-out electronics.

8.1.2 The timing system

In order to perform the feedback loops, to ensure the coherence between acquired data all over the site, a central timing is required to synchronise all Virgo subsystems. This is achieved with a master clock, driving a set of timing boards. The master clock is derived from a GPS system located in the Control Building [71]. Clock signals are distributed by optical fibre links to the different buildings (up to 3.2km away) where a local timing board converts them to TTL signals and VME bus interrupts for local use. The data, sampled at a frequency of up to 20 kHz, receive a time stamp with an accuracy of 1μs from these signals [72].

8.2 Global interferometer control

Virgo needs an active control of the positions of the suspended optical components, keeping the detector at its working point. The constraints are about 10^{-12} m rms for the longitudinal control ('Locking') and 10^{-9} rad rms for the angular degrees of freedom ('Alignment'). The dedicated hardware and software system, named Global Control [73], is in charge of the locking and the

2012 JINST 7 P03012

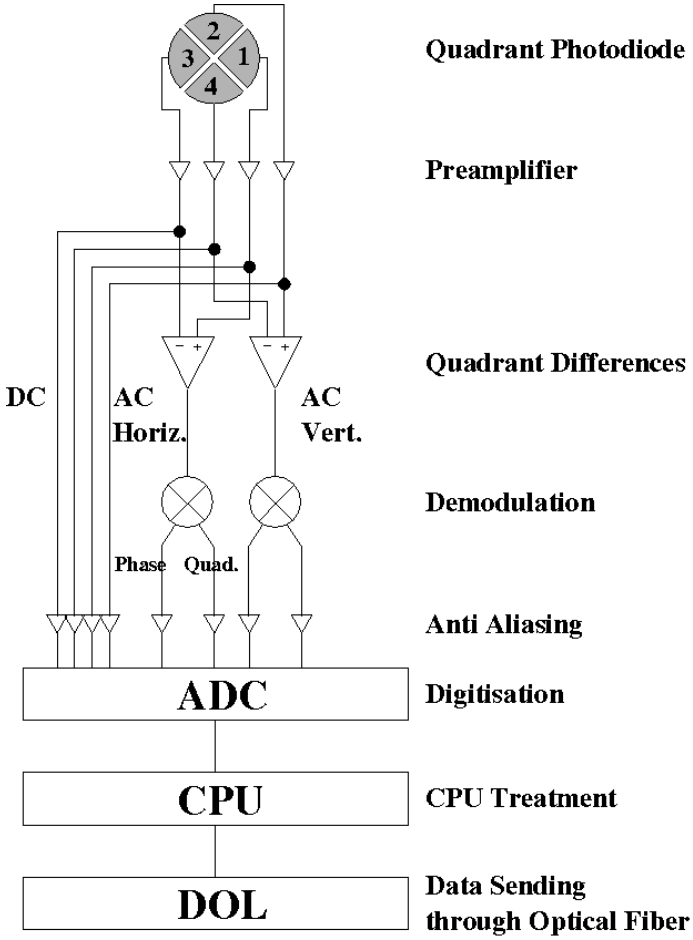

Figure 56. Quadrant read-out electronics.

alignment feedback loops. This system has been designed to provide a flexible tool for the integration of the various algorithms needed for the control of Virgo.

The locking feedback loop runs at 10 kHz and one of its major constraints is to be perfectly synchronised at this frequency. Any uncontrolled jitter or delay can introduce instability into the feedback loop. The Global Control receives, through the DOL, the digitised photodiode signals and computes the mirror position deviation from their optimal working point. The associated corrections are sent to each suspension control system to bring the mirrors back to the right position. The alignment feedback loop runs at 500 Hz and is similar to the Locking loop.

In order to fulfill the timing constraints and ensure a perfect synchronisation, two dedicated boards have been built. The first one, the Transparent Memory Board (TMB), allows for a fast access to the optical fibre boards in write and read mode. It also avoids the use of interrupts in the synchronisation, which are a source of uncontrolled delays. The second board, the Spy Data Board (SDB), is dedicated to the supervision of the processor running at 10 kHz. It gets a copy of all data exchanged at 10 kHz and delivers them to the monitoring machine.

The Global Control hardware is divided into three parts, as shown in figure 57: a crate housing the hardware dedicated to the locking feedback loops; a second crate for the alignment control loop; and a workstation monitoring the system.

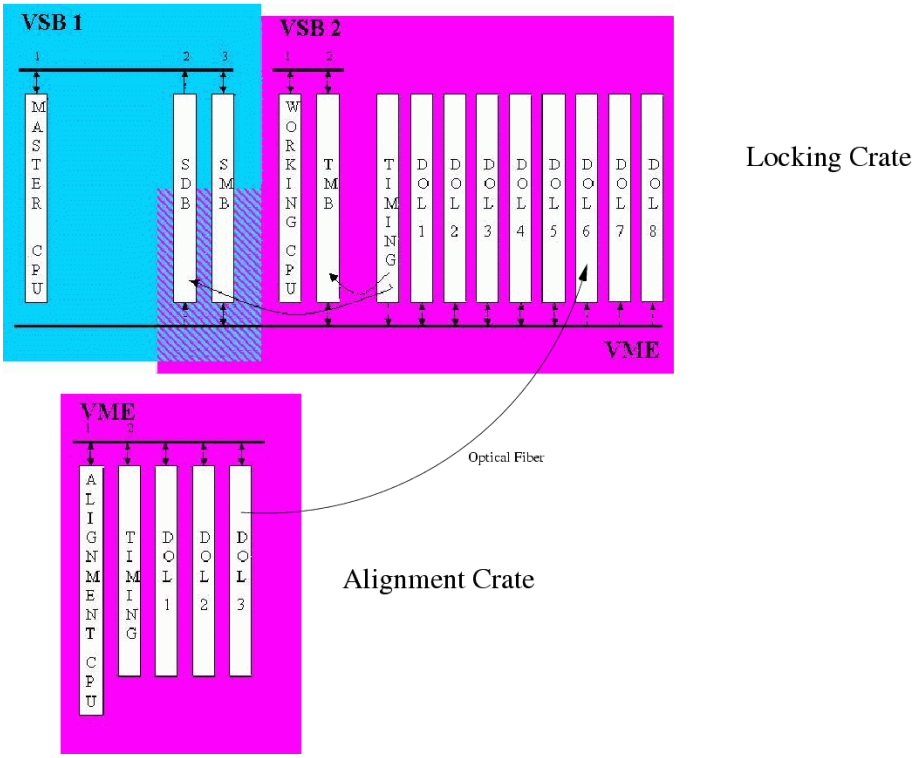

Figure 57. Schematic view of the Global Control crates.

The first crate is dedicated to the locking activities and the distribution of the correction signals. The second crate handles the alignment loop. In order to send the Alignment corrections to the suspensions, a DOL connects the locking and alignment crates. The two components of the locking part are the 'working' CPU, synchronised at the locking frequency, and the 'Master'. The latter steers the working part and transfers data for monitoring purposes. When synchronised at 10 kHz, the CPU housing the locking process is isolated from the network and communications go through a dual port VME/VSB Shared Memory Board (SMB).

The alignment CPU uses the VME bus to access the DOL boards. The synchronisation is performed using a VME interrupt generated by the timing board.

A general purpose CPU was chosen for the implementation of the control algorithms: this solution gives the best flexibility for the design of the feedback loops. In particular, the length reconstruction needs powerful algorithms requiring many mathematical operations that change during the lock acquisition process. The algorithms are split into three parts: sensing, filtering and driving. The Sensing part converts the photodiode signals into lengths and angles, the filtering part applies, in the time domain, the correction filters to the quantities computed by the sensing. In the driving part for locking, the various feedback signals are distributed to the relevant mirrors, according to the driving matrix given by the user. This algorithm division was chosen to allow switches between different versions of a given task without modifying the others. For example, it allows to change, on-the-fly, the length reconstruction procedure (from a lock acquisition to a lock maintaining algorithm, for example) while keeping the same filtering.

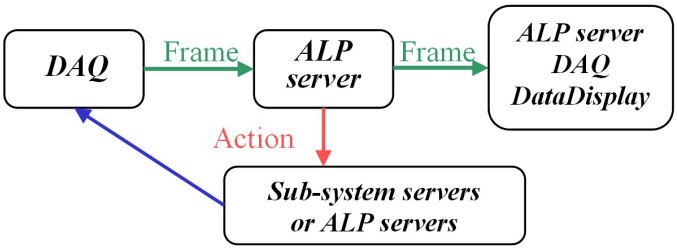

Figure 58. ALP simplified architecture

All free parameters of the algorithms (optical gains, filter characteristics, trigger thresholds, etc.) are stored in the Virgo on-line database. This allows the user to easily change the behaviour of the algorithms without modifying the C++ code. In addition, the full history of parameter changes is stored in the database, leading to an efficient recovery of previous configurations. Some parameters can also be modified on-the-fly without stopping the 10 kHz loop.

For both locking and alignment, the algorithm part of the code has been isolated in dedicated libraries which can also be used off-line on recorded data or connected to SIESTA [74], the simulation programme in the time domain developed for Virgo. This feature allows the complete development and test of an algorithm on an off-line workstation without perturbing the current version running in the on-line chain. In this case, the developer is not limited by real-time constraints and can use powerful tools to debug code. It should be underlined that the connection with SIESTA is a major improvement of the algorithm reliability, as the algorithm runs off-line in conditions as close as possible to the real ones. This feature really helped to improve the efficiency of the locking activities.

8.3 The Automatic Locking Procedure

The working point of the detector is reached and maintained using elaborate procedures, built on two levels:

- a real time and fast automation layer, directly implemented at subsystem level, involving local and global control, with a bandwidth up to 10 kHz;
- a control and monitoring automation layer, using the data coming from the DAQ channels to drive the different phases of the locking procedure.

The interaction between the two levels is implemented via the Virgo inter-process communication protocol Cm using DAQ channels.

The software developed for the control and monitoring level has a server architecture.

The basic concept is to use the data acquired by the DAQ and to compute on a standard workstation the state of any subsystem by processing the data from the relevant channels. According to the subsystem state, actions can be performed using system calls or directly with messages in Cm format (figure 58).

The ALP server uses the DAQ framework to access the data and to report on the performed actions. It uses a script programming language. It is possible, for example, to create internal

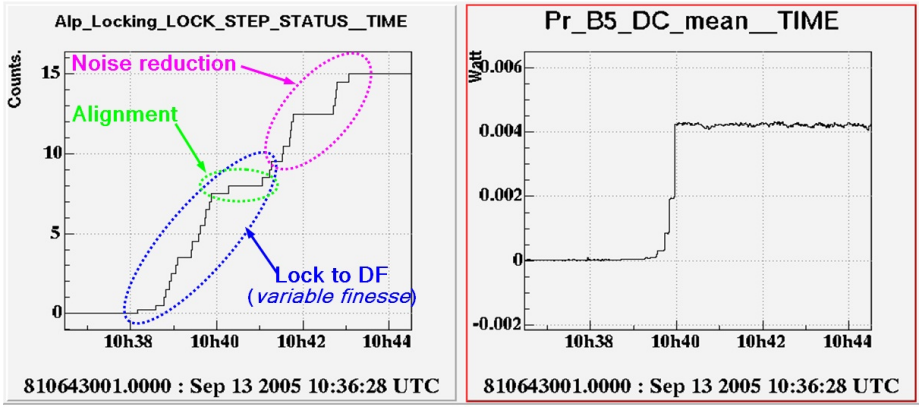

Figure 59. Left: locking status value transitions during the full sequence. Values correspond to the sequence step numbers. Right: power into the recycling cavity during the sequence evolution.

variables, compute signal statistical properties in the time domain (mean, min, max, rms, range) or in the frequency domain (FFT, rms in a frequency band, coherence, etc.) to perform the lock acquisition.

In order to better distribute the computational load and to provide a clean hierarchical organisation (subsystem-linked) to the different automation tasks, an automation architecture based on a Master and several ALP servers has been developed.

The overall automated procedure consists of two main sequences:

- Pre-alignment sequence. In this procedure the alignment of the beam in the North and West cavities is performed and the long optical cavities are tuned to resonance;
- Alignment/Locking/Noise Reduction sequence. This procedure allows to fully lock the interferometer along the longitudinal degree of freedom and introduces all controls for low noise operations.

  After having applied the main lock acquisition algorithm (called Variable Finesse [75]) it activates the angular drift control and the alignment control loop. It then tunes the Output Mode Cleaner. Next, it undertakes actions aimed at lowering the control noise and finally it adds permanent sinusoidal signals to the different mirror corrections for calibration.

During the sequence execution and after its completion, a special macro, working in the background, watches for any unlock event. In case of unlock, it brings the ITF to the lock acquisition pre-conditions, while avoiding further excitation of suspension motion.

This automation effort was an important element in achieving a duty-cycle higher than 80 % during data-taking. A typical locking sequence is shown in figure 59.

The risk of human error has been considerably reduced and daily commissioning activities, in which unlocking and re-locking the ITF is often required, have been greatly simplified.

2012 JINST 7 P03012

9 Environmental monitoring

The environmental monitoring system of Virgo records all of the physical parameters that can contribute or alter the response of the interferometer. In fact, due to the high sensitivity of the antenna and to the weakness of the expected signals, many external perturbations can influence the performances of the interferometer. The purpose of the environmental monitoring system is to generate vetoes against false gravitational wave signals in the data analysis phase. Such vetoes are based on the understanding of the noise generated by environmental factors.

The environmental monitoring system is composed of two main subsystems: the slow monitoring, also called Building Control, and the Fast Monitoring. The first subsystem [76, 77] records slowly changing parameters with a sampling frequency of 1 Hz. The second subsystem [77] records physical quantities with sampling frequencies of the order of some kHz. The actual sampling frequency depends on the characteristics of the probe being used. Actually, the slow control subsystem manages the acquisition of several outdoor and indoor parameters: weather conditions; lightning events; rain; temperature; pressure; and humidity.

Meteorological conditions are monitored using a commercial station (Vantage Pro model by Davis) equipped with an anemometer, a rain rate probe, a thermometer, a hygrometer and a barometer. These sensors are integrated in a single box placed on the roof of the Central Building of Virgo. The lightning detector is a commercial radio frequency detector, model LD-250 from Boltek, which is able to detect events, and recognise them as strikes, from a distance of up to 400 Km. This device is also positioned on the roof of the Central Building. Both the weather station and the lightning detector are connected, using a standard RS-232 link, to the VME-based master Building Control CPU.

The monitoring inside the Virgo buildings of the slow varying parameters, temperature, pressure and humidity, is performed using a highly modular and distributed acquisition system. The base element of this system consists of a commercial remote acquisition board, (TRS-56 model from Tecnint), equipped with 16 12-bit ADC channels. Conditioning electronic circuits adapt the signal range to the ADC input range. The standard transmission protocol used, RS-485, makes it possible to connect several boards in a cascading style over a distance of 300 metres. The master board is a TVM-932A VME board from Tecnint, which resides in a crate together with a main CPU collecting all of the data and sending them to the acquisition system. Due to the slow acquisition frequency (1 Hz), no special synchronisation is required to manage these processes. There is one such system serving the Central and Mode Cleaner Buildings and two other systems for each end building. The total number of slow monitoring probes actually operational in Virgo is: 244 temperature, 11 pressure and 11 humidity probes, 1 weather monitoring station and 1 lightning detector station.

The Fast Monitoring subsystem records physical quantities with a broader frequency band. The system is composed of mechanical vibration sensors, acoustic sensors and magnetic field sensors, as listed in table 7.

The location of the various sensors is shown in figures 60 and 61.

In particular, vibration and acoustic sensors are located near the Superattenuator base and on the optical benches. The mechanical vibration sensors are of three kinds. The first consists of single axis accelerometers, 4378 model from Brüel and Kjaer, with a frequency band spanning from 1 Hz

2012 JINST 7 P03012

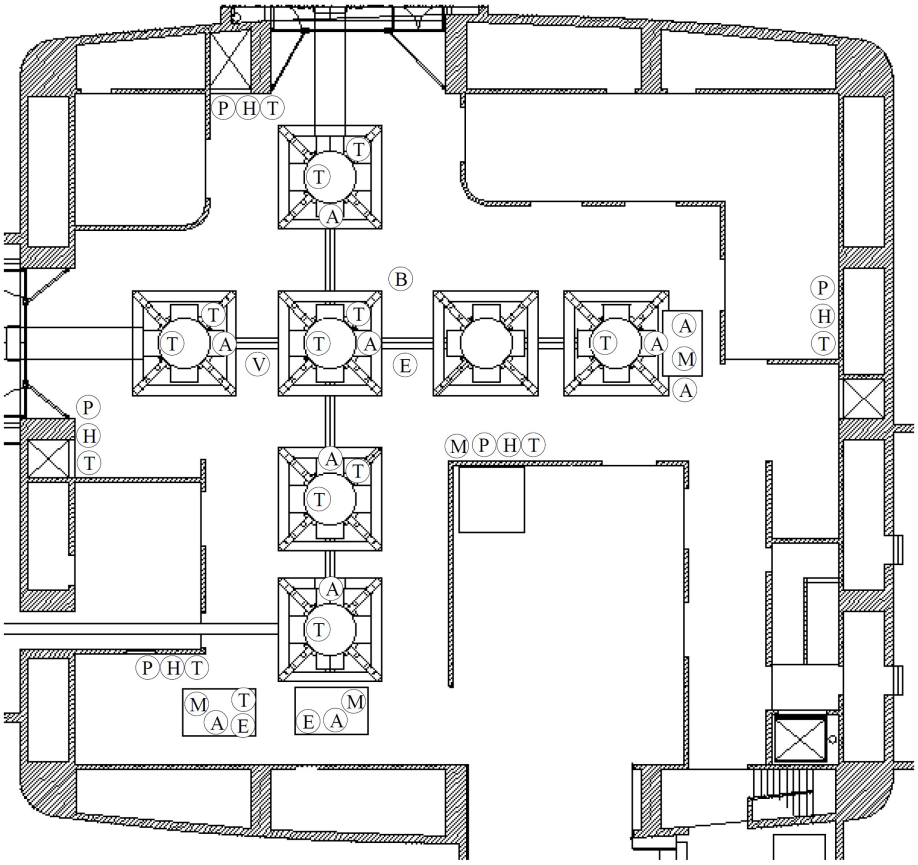

Figure 60. Location of the various sensors in the Central Building. M: Microphone, B: Magnetometer, E: Episensor, V: Velocimeter, A: Accelerometer, T: Temperature probe, P: Pressure probe, H: Humidity probe.

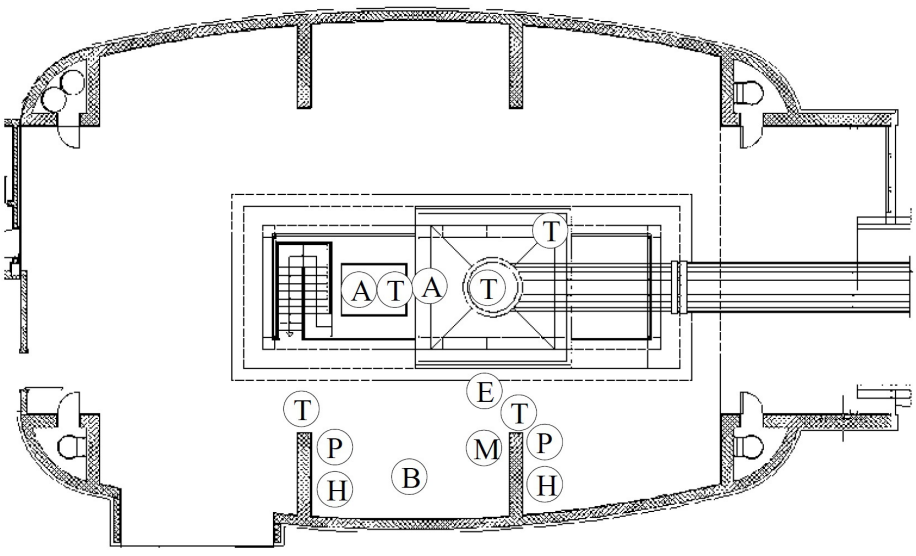

Figure 61. Location of the various sensors in the End Buildings. Same labels as used in figure 60.

Table 7. List of the sensors used in the environmental monitoring system.

Environmental sensors			
Model	Physical quantity	Number of axes	Frequency band (Hz)
Kinemetrics MFS-06	Acceleration	3	0.1 – 200
Brüel and Kjaer 4378	Acceleration	1	0.1 – 3000
Güralp CMG-40T	Velocity	3	0.03 - 50
Brüel and Kjaer 4190	Pressure (Acoustic)	1	0.1 - 3000
Metronics MFS-06	Magnetic Field	1	4 - 10000

to 3 kHz. The second kind includes tri-axial accelerometers, working in the frequency range 0.1 to 200 Hz, the FBA ES-T model from Kinemetrics. The last type is a tri-axial broadband velocimeter, CMG-40T model from Güralp, the frequency band of which spans from 30 mHz up to 50 Hz. The acoustic sensors are microphones, model 4190 from Brüel and Kjaer, with a frequency band of 0.1 to 3 kHz. Finally, the magnetic sensors are broadband induction coil magnetometers, MFS-06 model from Metronix, with a frequency band from 4 Hz to 10 kHz.

All the fast monitoring probes are connected to the ADC of the data acquisition system of Virgo [section 10]. For microphones and single axis accelerometers, a commercial amplifier, a Nexus model from Brüel and Kjaer, is used, while bespoke conditioning amplifiers are used for the remaining sensors.

The large number of probes used for the monitoring and their optimal positioning makes it possible to understand several properties of the environmental noise on the site and to study its coupling with the Virgo interferometer [78, 79].

2012 JINST 7 P03012

10 Data acquisition and storage

10.1 Data acquisition (DAQ)

The task of the data acquisition system [80] is to collect, in a coherent way, all of the information required for the validation of detector data and gravitational wave signal extraction. Therefore, it requires a site-wide timing system to produce and synchronously record all of the data; different dedicated read-out systems or interfaces to the control servos distributed over the site; tools to format, transport and write the data in files; and on-line monitoring tools to display the data or to analyse in real-time the state of the detector.

The DAQ is a fundamental component of the Virgo control system, which is digital for a very large fraction. The timing system is the same as that used to drive the controls (section 8). The digitised signals used by the servo loops are collected by the DAQ. This is achieved without introducing any disturbance or delay into the control loops.

One important requirement relating to the DAQ system is its need to be flexible and scalable, in order to accommodate the evolution of the interferometer during the commissioning phase. In addition, reliable continuous operation must be ensured during the whole life-time of the experiment, while following technological progress.

10.1.1 Data acquisition choices

The high number of channels to be read-out in Virgo requires a DAQ system that is able to sustain several MBytes of data flux and a high data storage capacity with fast random access to the data. Further requirements are imposed by the distribution of the DAQ over the two 3 km arms. In addition, on-line control constraints require a latency of at most a few seconds between ADC read-out and availability of data for display in the Control Room. Finally, the production of several monitoring and trend data streams is needed to follow the state of the detector and to monitor the performance of the data acquisition system.

The number of channels has evolved according to the comprehension of the detector behaviour and the phases of the commissioning. The necessary flexibility has been achieved using modular hardware and software building blocks. Standard Unix-like workstations with a set of libraries and 'plug-and-play-like' software tools compose the data acquisition system. Shared memories and Ethernet connections are used for communication. Some Virgo in-house developed basic libraries [81] are used to easily handle in a uniform way the configuration management, the on-line monitoring, the data transfer or the on-line access to the data. The same data format, called the 'frame' format [82], is used for data acquisition and on-line processing. This provides an easy way to store data as time slices of various lengths and to produce different data streams. This frame format has been adopted for data exchange between the various gravitational wave detectors around the world.

This modular data acquisition system has a three-level structure: the front-end read-out, the frame builder and the frame collector.

10.1.2 The front-end read-out

The front-end read-out processes are synchronised over the whole site using a timing system (section 8) based on a GPS signal with micro-second precision. Real time processors (RIO 8062 PowerPC) running under LynxOS manage the exchange of data with the electronics modules hosted in

2012 JINST 7 P03012

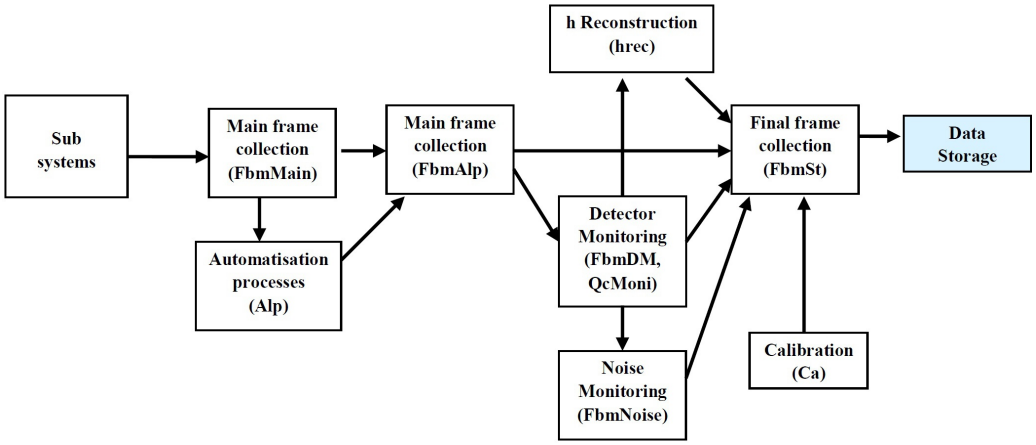

Figure 62. Frame builder architecture.

VME crates. These include ADC, timing boards, DSP, DAC and digital camera read-out. Time critical data transmission is achieved using Digital Optical Links (DOL), while other data are exchanged with an Ethernet connection.

10.1.3 The frame builders

These elements collect the data from the front-end read-out and put them in frame format:

- The Slow Frame Builder: collects the data sent through Ethernet by the Slow Monitoring Stations, which provide data at a rate of less than or equal to 1 Hz. Periodically, the timing system sends a message to the slow frame builder, the slow frame builder then requests data to the various Slow Monitoring Stations through Ethernet. Data are sent back as an ASCII string and are put in frame format. The slow frame builder periodically sends these frames to the frame collector;
- The Fast Frame Builder: collects data sent through optical links by the front-end processors for photodiodes and suspensions sensors. The data are stored in the buffer of the reception DOL. The packets collected from various DOL and stamped with the same frame number are merged into one second frames. A data compression algorithm reduces the size of these frames by a factor of about 2.5 in stable interferometer conditions;
- Specific Frame Builders: some processes, in charge of specific actions (detection bench control, output mode cleaner control, camera read-out or the environmental monitoring) read ADC when receiving an interrupt signal provided by the timing system. They then format the data into frames and send them directly to a frame collector.

The Frame Builder architecture is shown in figure 62.

10.1.4 The frame collectors

Also named Main Frame Builders, these asynchronously collect the frames sent by the various frame builders and merge them into a single frame containing all of the acquired data. These data

2012 JINST 7 P03012

are then dispatched to the next stage of the DAQ pipeline using a shared memory. This mechanism allows for a fast and easy on-line access to the data. Several 'consumer' processes can read this shared memory, each with its own channel selection, its own processing and its own output destination. All the frame collectors run on standard workstations under a Unix-like operating system.

Each main Virgo subsystem (suspensions, detection, injection and environmental monitoring) has its own data acquisition line with the three data acquisition levels described above. This allows, when needed, an independent on-line access to the data of each subsystem through its frame collector. It also makes it possible to distribute first-level on-line processing tasks that consume much CPU power. At the output of the four data acquisition lines a final frame collector sends frames to the raw data storage and to the on-line workstations.

10.1.5 Raw data flow and latency

At the output of raw data collection, each frame (one second of data) contains about 7 MBytes of compressed data (about 18 MBytes per second if uncompressed). Fast Ethernet and Gigabit Ethernet networks with TCP/IP protocol are used to transmit these data down to the on-line display programmes, to monitoring processes and to the data storage. At this level, the data are available with a delay lower than 2 seconds.

10.1.6 Data streams

To complement the raw data stream as well as to provide summary information over long time segments, trend data with various sampling frequencies are produced. These are the streams of the 50Hz re-sampled data and one second and one minute trend data that also contain the results of simple statistical analyses. These streams are produced on-line and stored on-disk, providing information that now spans several years.

An additional data stream, containing spectra of selected channels, is produced in order to have fast access to information in the frequency domain (for instance to perform time-frequency plots over hours or days).

10.1.7 DAQ system monitoring

A graphical interface has been developed to control all DAQ processes. This graphical interface uses the Ethernet network and the Virgo basic libraries to provide both in the Control Room and on a web page the status of the DAQ processes (started or stopped, memory consumption, etc.). In addition to this interface, some trend data are produced within the DAQ chain to perform a statistical analysis of the DAQ system performance.

10.1.8 DAQ system performance

The main characteristics and performance of the DAQ system are the following:

- overall mean latency: 1.8 sec;
- mean data flux: 30 MB/s (uncompressed) 8-9 MB/s (compressed);
- mean number of raw data channels: 1350;

- mean number of trend data channels: 7600;
- frame length: 1 s;
- trend frame length: 1800 s;
- GPS time stamp precision: 1 μs.

The channel sampling rate varies from 20 kHz to less than 1 Hz depending on the bandwidth of the measured physical quantity.

10.2 On-line monitoring

10.2.1 Data quality and detector monitoring

A set of monitoring programmes provide information (mainly quality flags) tagging the time epochs when data are not use-able for analysis. This information is also provided on-line in the Control Room and on web pages that show the status of the complete interferometer and its subsystems, and suggest possible operator actions. A summary quality flag represents the overall detector status and can be used to prevent analysis in inappropriate conditions. In addition, a 'Science Mode' flag is generated upon the request of the operator, when a basic set of quality checks are fulfilled. These flags are recorded in the data frames so that they can be used in the data analysis as a first-level veto.

10.2.2 Data Display

The Data Display is a programme developed for-purpose to provide an easy way to display data on-line or off-line and to perform some minimal data analysis. It acts as a digital oscilloscope and spectrum analyser for the signals present in the interferometer. Histograms of signal amplitudes and basic statistical analysis are available. Channels can be combined to study coherence, correlation or to measure transfer functions. The Data Display can also be run remotely or on off-line data and is a widely-used tool in interferometer studies.

10.2.3 Web monitoring

Relevant trend data are displayed on web pages using a tool that provides periodically updated plots for the main channels of each detector subsystem, for the data acquisition or for on-line processing and data analysis: http://wwwcascina.virgo.infn.it/MonitoringWeb.

10.3 Storage and data management system

Data storage has to comply with a continuous data flow of 8-9 MB/s, while being highly reliable, in order to guarantee the recording of rare astronomical events. In addition, due to the continuous tuning work undertaken on the detector, the most recent whole set of raw data (about 5 months) needs to be available for analysis on-site. This results in the need for a total of 200TB of disk storage farm, 146 of them reserved for the raw data circular buffer. The data storage must allow for the writing of data by the DAQ system, simultaneous access by users and continuous transfer to the computing centres (CNAF-Bologna and CCIN2P3-Lyon) which act as external data repositories. For archive and safety, data are also replicated on tapes.

2012 JINST 7 P03012

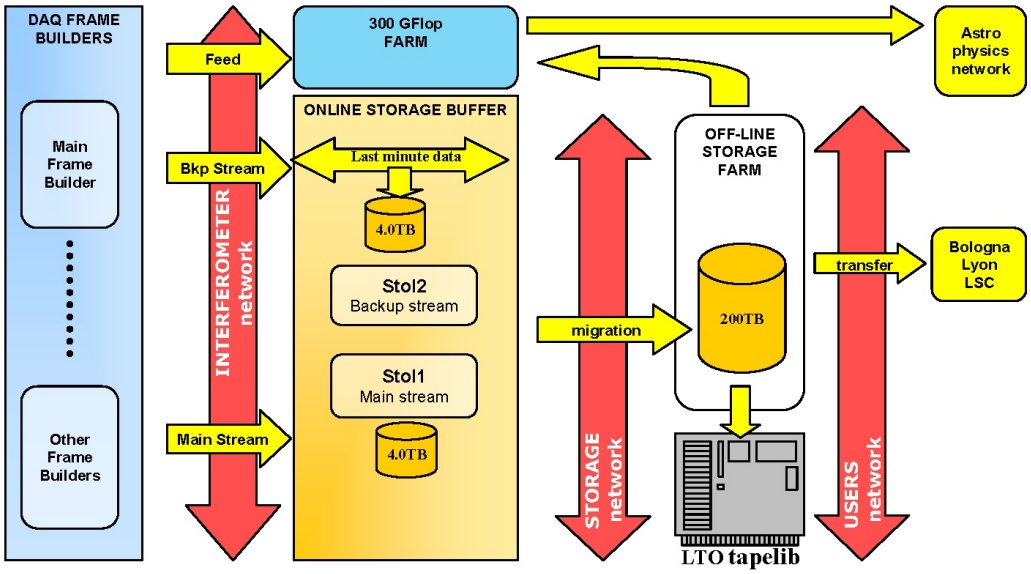

Figure 63. Storage architecture.

Data are written using dedicated systems of workstations and on-line buffers, while main user access is through the storage farm (off-line storage). The on-line and off-line storage systems are connected by a dedicated LAN and data are migrated asynchronously from the on-line buffers to the storage farm. From the storage farm, data are then sent to the external computing centres, where most of the analysis is performed. Figure 63 shows the general storage architecture.

10.3.1 On-line storage

For the raw data, i.e. nearly the total of the produced data, two identical streams are spilled from the Main Frame Builder and sent to two on-line buffers ('stol1' and 'stol2') over Ethernet connections. The buffers are two independent Linux servers, where the DAQ clients write on directly connected storage arrays.

'Stol1' is considered as the reference stream, therefore the storage is a fully redundant 4 TB FC-to-FC RAID5 high-performance array with hot spare disks. It guarantees data buffering for about 1 week in case of failure of the downstream storage.

Data in 'stol1' are not directly accessible by users, but become available once they have been migrated to the storage farm.

'Stol2' is an identical server but with an FC-to-SCSI array of 4TB, its main purpose being to provide a redundant copy of the data in case *stol1* has some faults.

The 'stol2' server exports data directly to the user network via NFS, giving a direct access to the last-minute data.

Data sampled at 50Hz and trend data are written on another non-redundant server ('stol3') with 2TB of local disk space. This kind of data is replicated to a dedicated off-line storage area for archiving. These archives are also sent to the remote repositories.

10.3.2 Off-line storage

The storage farm is composed of 10 Linux servers connected to a 200 TB Storage Area Network, with a mixed technology and with a maximum read/write speed of 2 GB/s. The disk area is seen by the local users as NFS volumes of 1 TB each that can be accessed at the speed of 40 MB/s.

A pool of back-up clients runs directly on the storage farm servers and save the raw data files to a 200TB tape library with LTO Ultrium2 tapes. The tapes are vaulted at a rate of 10 tapes per week and put on shelves for crash recovery of the storage farm volumes.

The tape library uses 3 tape drives even if 1 drive is nominally capable of sustaining a data transfer rate of 80MB/s.

10.3.3 Computing farm

The in-time Data Analysis algorithms for triggers and noise monitoring run on a 96 node computing farm with a total of 320 CPU cores, with aggregated capacity of more than 300 GFlops. The nodes are based on the AMD Opteron processor architecture, which proved effective in the SIMD (single instruction, multiple data) computation common to many GW data analyses [83].

The architecture has been optimised for direct on-the-fly feeding from the frame builders via TCP connections and distribution on the internal gigabit interconnect network, but data on the storage farm can also be accessed, allowing computations over longer time windows.

All the nodes are functionally independent to allow different analysis pipelines to run without a centralised batch scheduling system; only the Scientific Linux base platform configuration is managed centrally.

10.3.4 Data management

The absence of a single virtual file system where data can be placed, irrespective of the physical storage and the possibility of hardware failures, requires a special software to move the data in a safe way from the on-line buffers to the storage farm (where most recent data are stored in a circular buffer) keeping track of the file positions and replicas.

The storage nodes keep track of local files, while a global placement map of migration status and volume availability is maintained. The information relating to file location and frame content is written in a Frames File List (FFL), which is the common format used by the software applications to locate the data.

The software also selects the best replica available on the on-line storage, comparing the two raw data flows, then moves the data to the off-line buffers (deleting the oldest files when needed) and schedules them for the tape archiving and for the data transfer to the remote repositories.

The data archive (the entire off-line storage) is exported to the user network via an NFS directory structure handled by an auto-mounter daemon.

10.3.5 Data transfer

The data transfer architecture involves two dedicated servers (‘datagw’ and ‘dataldr’) that see the off-line storage farm as a set of NFS volumes.

The ‘dataldr’ machine is remotely controlled and handles the exchange of data with the LIGO Scientific Collaboration using the LDR software, reading files directly from the off-line storage and saving the incoming data into a dedicated buffer (7TB).

2012 JINST 7 P03012

The ‘datagw’ machine takes care of the replica to the remote repositories in the national computing centres of Lyon and Bologna; it hosts some software procedures, integrated with the data management software, that handle the replica to process by means of two different engines (‘bbftp’ for Bologna and ‘SRB’ for Lyon) which match the different architectures available locally at repositories [84].

It also produces a detailed log for all operations performed and the metadata file (the Frame File List) for both destinations, based upon information stored in a local MySQL database.

11 The Virgo effective sensitivity and on-line strain computation

This last chapter shows how the Virgo spectral sensitivity to gravitational wave signals is obtained from the apparatus described in this paper; this needs a precise calibration of the detector considered as a whole, once each specific subsystem has been characterised.

Using the important on-line computing power and data storage available on site, this sensitivity can be computed nearly in real-time during the running periods, making it possible to check the noise status of the instrument and to act quickly if problems are detected.

The instantaneous strain output of the detector, called 'h reconstructed', is computed from the dark fringe signal and the calibration. This is the main output of Virgo; it is computed and written on each frame before final storage.

11.1 Detector calibration

The precise recording of a gravitational wave signal based on the relative elongation of the arms requires an accurate calibration of the interferometer response to mirror motion.

Since the ITF is controlled, mirror motion is reduced compared to that which should have been induced by the gravitational wave. It is thus also necessary to take into account the mirror displacement due to the control signals. Therefore, the calibration consists of two tasks:

- to measure mirror displacement due to a given control signal in order to obtain the transfer function for each mirror actuator;
- to measure the effect of a differential arm elongation on the dark fringe signal to compute the detector response to gravitational waves.

11.1.1 Mirror actuator transfer function

The procedure to measure the mirror actuator transfer function (TF) is based on the distance between interferometer fringes, determined by the laser wavelength. Due to the high force required this can only be performed using coil drivers in High Power (HP) mode. The Low Noise (LN) channel of the coil driver is then calibrated electronically.

First, data are taken with the ITF in free-swinging Michelson configurations. Sinusoidal signals (zCorr) are sent to the mirror actuators with the coil driver set in HP mode. In such a configuration, it is possible to extract the mirror movement ΔL from the un-calibrated dark fringe signal directly. The ratio ΔL/zCorr is the mirror actuator TF in HP mode at each line frequency measured in m/V. Such measurements are done from 5 Hz to 1 kHz on the beam splitter and arm mirrors.

To convert the actuator TF from HP to LN mode, white noise is sent to the actuator. The current flowing in the actuator coil is measured as a voltage across a resistor with the coil driver in HP and LN mode. This makes it possible to compute the mirror actuator TF in LN mode.

Since the mirror position is partially controlled by the marionette, a measurement of the related TF is required. Measurements in the free-swinging Michelson configurations are not sensitive enough above a few Hz. The measurement is thus done with the fully locked ITF. The ITF response to an unknown differential arm elongation (10Hz – 100Hz band) sent through the marionette actuation is compared to the response to a known differential arm elongation sent through the calibrated mirror actuation [85].

2012 JINST 7 P03012

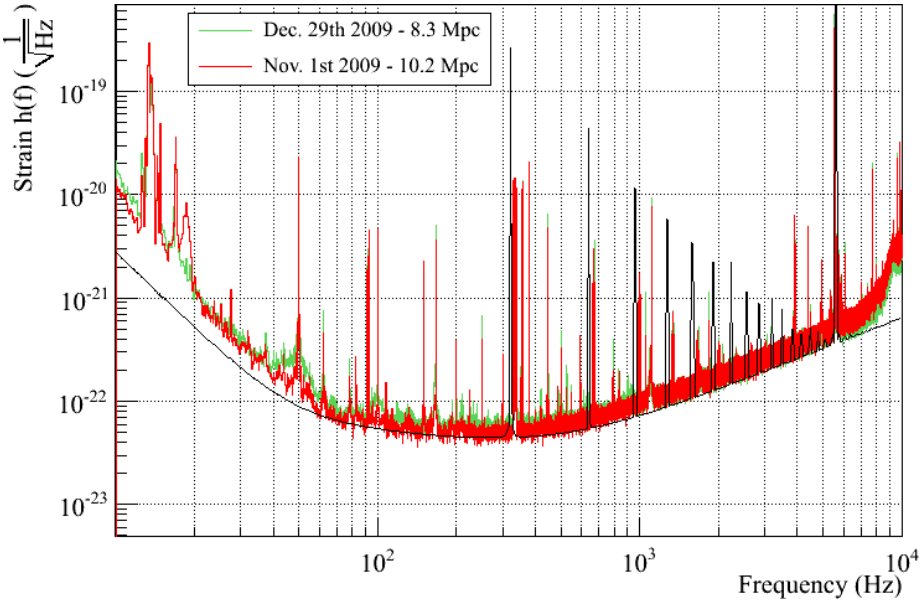

Figure 64. Virgo sensitivity in November 2009 (red), December 2009 (green) and design curve (black).

11.1.2 Virgo transfer function and sensitivity

The Virgo transfer function describes the effect of a differential arm elongation in metres on the dark fringe intensity in Watt, in the range 10 Hz to 10 kHz. It is measured by injecting white noise from 10 Hz to 1 kHz on the NE or WE mirror actuator and measuring the dark fringe signal. From the mirror actuator transfer function, the injected white noise in Volt can be converted into a mirror displacement in metres ΔL. The ratio between dark fringe intensity variation and ΔL is a measurement of the Virgo TF up to 1 kHz. Above that frequency the Virgo TF shape is determined by the 3 km cavity pole behaviour. The pole frequencies have been measured to be at 500 Hz within a few per cent.

The displacement sensitivity in m/Hz$^{1/2}$ is obtained by multiplying the dark fringe noise linear power spectrum in W/Hz$^{1/2}$ by the Virgo TF in m/W. The strain is obtained by dividing the result by the arm length. The best sensitivity, obtained during November and December 2009, is shown in figure 64, together with the design sensitivity.

11.2 On-Line Preliminary Analysis

The on-line computing capability on-site has been designed to make it possible for several tasks to be performed on the data as soon as they are acquired. Different algorithms have been written, some of them for data conditioning, such as down sampling. The crucial task of h-reconstruction is also performed, meaning it is possible to immediately check the instrument sensitivity and to identify anomalous noise. To monitor the time stability of the detector sensitivity for frequencies between 70 Hz and 1300 Hz (frequency of the last stable orbit for inspiralling neutron stars) the 'horizon' is computed on-line. This quantity is defined as the largest distance at which a typical binary neutron star inspiral event can be detected by Virgo with a signal-to-noise ratio equal to 8. An average is taken over all possible orbit orientations and source positions in the sky. The inputs

2012 JINST 7 P03012

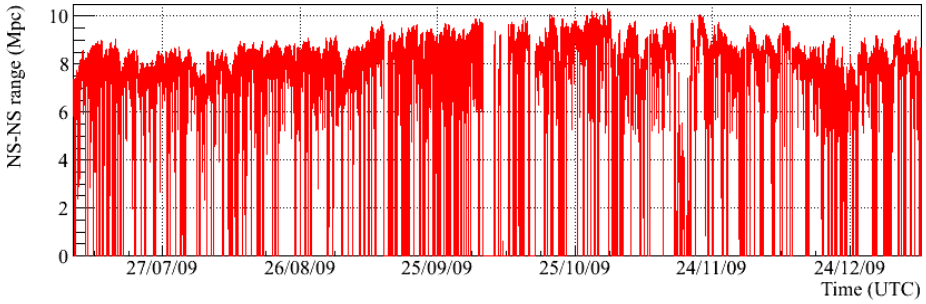

Figure 65. Binary neutron star averaged horizon (range) from August to December 2009. The variations are mainly due to weather conditions and control noises.

for this calculation are the noise spectrum and the expected waveform [86]. Figure 65 shows this averaged horizon (range) during the second part of 2009 (August to December).

Finally, algorithms are being designed to detect on-line hints of possible event candidates by looking for transient signals from expected sources, such as inspiralling binary systems and stellar core collapses (details of possible methods can be found in [87, 88]). It is planned to send these signals as triggers to telescopes and satellites observing the whole electromagnetic spectrum and to neutrino observatories. Agreements with several detectors are already operative.

The final data analyses, which require a huge computing power, are performed off-line in the large computing centres.

2012 JINST 7 P03012

12 Conclusion

Given the very stringent and mandatory technical constraints, the creation of the Virgo interferometer was a real challenge. Nevertheless, all of these difficult constraints have been met. The longest and most difficult task was to control the long cavity length to better than 10^{-15} metre (rms), mirror angles to better than 10^{-9} radian, while achieving a mirror position noise of 2 10^{-19} m $\mathrm{Hz}^{-1/2}$ at 150 Hz.

The present achievement of the Virgo detector means a sensitivity very close to the design for frequencies above 100 Hz, putting Virgo among the best GW detectors. Below 40 Hz, Virgo is by far the most sensitive detector, almost at its design sensitivity. The observed excess noise has been identified to be generated by diffused light and actuator noise; suitable solutions are being prepared. Reliability has also been demonstrated and several months of data for GW searches have already been collected.

Nevertheless, even if the present sensitivity represents enormous progress with respect to previous detectors, the latest astrophysical predictions for the detection rate are still low. However, measuring wave amplitude, rather than energy, the former decreasing as *1/R* , makes it possible to profit even from the smallest improvement in sensitivity. A reduction of a factor 2 in the noise floor allows for the exploration of a volume of universe 8 times larger, increasing the detectable source population by about the same factor. An upgrade programme is the obvious consequence of these considerations.

The commissioning of the Virgo detector has shown that the physical principles underlying it are well understood and that, following the removal of remaining technical noises, the sensitivity is going to be determined by the fundamental noise sources: shot noise, thermal and seismic noise. By having chosen to use Superattenuators, Virgo will be able to decrease seismic noise down to the limit imposed by fluctuations of the gravity gradient at low frequency. Further improvements in that direction would probably require locating the detector underground. On the other hand, upgrades of the detector can reduce shot noise, by increasing the power of the light source, and thermal noise, by using different materials, like silica, for the mirror suspensions.

The first Virgo upgrade (called Virgo+) includes the increase of the laser power up to above 50Watts, coupled with the introduction of a system of thermal compensation to homogenise the temperature distribution on the large input mirrors; this is completed with an infrared radiation curvature correction system. A second improvement is the use of silica fibres to suspend the mirrors. All this seems to be quite efficient in the reduction of thermal noise, improving the sensitivity especially at low frequency (20-60 Hz). Finally, increasing the Finesse of the cavities from 50 to 150 completes this first step. The detector has been run with these modifications during the past few months, and they will all be detailed in a forthcoming paper, together with an analysis of results.

The second upgrade step, known as Advanced Virgo, will include major changes, mainly in the optical scheme, introducing a signal recycling mirror and increasing the laser power up to 200 W. The foreseen changes should result in an increase of the sensitivity of one order of magnitude in the whole detection band. The accessible universe volume should then be increased by a factor of 1000, compared to the Virgo design. We would then expect a detection rate of at least a few events per year.

2012 JINST 7 P03012

References

[1] A. Einstein, *Näherungsweise Integration der Feldgleichungen der Gravitation,* Sitzungsberichte der Preußische Akademie der Wissenschaften **1** (1916) 688.

[2] R.A. Hulse and J.H. Taylor, *Discovery of a pulsar in a binary system, Astrophys. J.* **195** (1975) 51.

[3] J. Weber, *Detection and Generation of Gravitational Waves*, *Phys. Rev.* **117** (1960) 306.

[4] E. Amaldi et al., *First gravity wave coincidence experiment between resonant cryogenic detectors — Louisiana-Rome-Stanford*, *Astron. Astrophys.*, **216** (1989) 325;
I. Heng, D. Blair, E. Ivanov and M. Tobar, *Long term operation of a niobium resonant bar gravitational wave antenna*, *Phys. Lett.* **A 218** (1996) 190;
L. Baggio et al., *3-mode detection for widening the bandwidth of resonant gravitational wave detectors*, *Phys. Rev. Lett.* **94** (2005) 241101 [gr-qc/0502101].

[5] C. Bradaschia et al.: *VIRGO: proposal for the construction of a large interferometric detector of Gravitational waves*, CNRS-INFN, unpublished (1989).

[6] B. Abbott et al., *Detector description and performance for the first coincidence observations between LIGO and GEO*, *Nucl. Instr. Meth. Phys. Res.* **A 517** (2004) 154.

[7] LIGO SCIENTIFIC collaboration, B. Willke, *GEO600: Status and plans*, *Class. Quant. Grav.* **24** (2007) S389.

[8] R. Takahashi et al., *Operational status of TAMA300 with the seismic attenuation system (SAS)*, *Class. Quant. Grav.* **25** (2008) 114036.

[9] D. Blair et al., *The Science benefits and preliminary design of the southern hemisphere gravitational wave detector AIGO*, *J. Phys. Conf. Ser.* **122** (2008) 012001

[10] H. Araujo et al., *LISA and LISA PathFinder, the endeavour to detect low frequency GWs*, *J. Phys. Conf. Ser.* **66** (2007) 012003 [gr-qc/0612152] [INSPIRE].

[11] R.W.P. Drever, *Gravitational Radiation,* Les Houches Summer Institute (1982).

[12] G. Rupschus, R. Niepraschk, K. Jousten, and M. Kühne, *Detection of small pressure pulses in an ion pumped ultrahigh vacuum system*, *J. Vac. Sci. Technol.*, **A 12** (1994) 1686.

[13] S. Braccini et al., *Test of Virgo prototype vacuum system*, *Vuoto* **23** (1994) 41.

[14] P. Marin, M. Dialinas, G. Lissillour, A. Marraud and A. Reboux, *Outgassing performances of an industrial prototype tube for the Virgo antenna*, *Vacuum* **49** (1998) 309.

[15] V. Brisson, J. Hoang, G. Lissilour, P. Marin, A. Reboux, M. Bernardini and A. Pasqualetti, *Ultra-high vacuum qualification of the prototype module for the 2×3 km arms of the VIRGO interferometer*, *Vacuum* **60** (2001) 9.

[16] M. Bernardini et al., *Air bake-out to reduce hydrogen outgassing from stainless steel*, *J. Vac. Sci. Technol.*, **A16** (1998) 188.

[17] M. Dialinas, *Virgo Note*, unpublished, VIR-NOT-LAL-1390-083, (1997).

[18] M. Bernardini et al., *Virgo Note*, unpublished, VIR-SPE-PIS-3100-106 (1998).

[19] A. Reboux, *Virgo Note*, unpublished, VIR-MEM-LAL-3100-166, (1999).

[20] V. Dattilo, *Virgo Note*, unpublished, VIR-TRE-PIS-4600-145 (1998).

[21] J.L. Beney et al., *Virgo Note*, unpublished, VIR-TRE-PIS-3400-159 (2000).

2012 JINST 7 P03012

[22] K.S. Thorne, *Light scattering and proposed baffle configuration for the LIGO, Caltech Theoretical Astrophysics Preprint* N. GRP-200 (1989).

[23] J-Y.Vinet et al., *Scattered light noise in gravitational wave interferometric detectors: A statistical approach*, *Phys. Rev.* **D 56** (1997) 6085.

[24] J.-Y. Vinet, V. Brisson and S. Braccini, *Scattered light noise in gravitational wave interferometric detectors: Coherent effects*, *Phys. Rev.* **D 54** (1996) 1276.

[25] S. Braccini, *Seismic Isolation in VIRGO*, Ph.D. thesis, Scuola Normale Superiore, Pisa, (1996) unpublished.

[26] G. Ballardin et al., *Measurement of the VIRGO superattenuator performance for seismic noise suppression*, *Rev. Sci. Instrum.* **72** (2001) 3643.

[27] VIRGO collaboration, *Virgo Final Design*, E.T.S., Pisa, Italy (1995).

[28] G. Losurdo et al., *An inverted pendulum pre-isolator stage for the Virgo suspension system*, *Rev. Sci. Instr.* **70** (1999) 2507.

[29] G. Losurdo, G. Calamai, E. Cuoco, L. Fabbroni, G. Guidi, et al., *Inertial control of the mirror suspensions of the VIRGO interferometer for gravitational wave detection*, *Rev. Sci. Instrum.* **72** (2001) 3653 [gr-qc/0105110].

[30] M. Beccaria et al., *Extending the Virgo gravitational wave detection band down to a few Hz: metal blade springs and magnetic antisprings*, *Nucl. Instr. Meth. Phys. Res.* **A 394** (1997) 397.

[31] M. Beccaria et al., *The creep problem in the Virgo suspensions: a possible solution using Maraging steel*, *Nucl. Instr. Meth. Phys. Res.* **A 404** (1998) 455.

[32] S. Braccini et al., *The maraging-steel blades of the Virgo super attenuator*, *Meas. Sci. Technol.* **11** (2000) 467 S. Braccini et al., *The maraging-steel blades of the Virgo super attenuator*, *Measur. Sci. Tech.* **11** (2000) 467 [*Erratum ibid.* **15** (2004) 599].

[33] A. Delapierre and F. Frasconi, *Stress distribution on the suspension wires of the Virgo superattenuator*, Virgo Internal Report VIR-TRE-PIS-4600-137, (1997).

[34] S. Braccini et al., *An improvement in the Virgo Super Attenuator for interferometric detection of gravitational waves: The use of a magnetic antispring*, *Rev. Sci. Instr.* **64** (1993) 310.

[35] G. Ballardin et al., *Measurement of the transfer function of the steering filter of the Virgo super attenuator suspension*, *Rev. Sci. Instrum.* **72** (2001) 3635 [*Erratum ibid.* **73** (2002) 233].

[36] G. Cagnoli, L. Gammaitoni, J. Kovalik, F. Marchesoni and M. Punturo, *Full scale prototype of high Q pendulum for interferometric gravitational wave detectors*, *Rev. Sci. Instrum.* **71** (2000) 2206.

[37] P.R. Saulson, *Thermal noise in mechanical experiments*, *Phys. Rev.* **D 42** (1990) 2437 ;
S. Rowan et al., *Q factor measurements on prototype fused quartz pendulum suspensions for use in gravitational wave detectors*, *Phys. Lett.* **A 233** (1997) 303.

[38] M. Punturo, *The Virgo sensitivity curve*, Virgo Internal Report No. VIR–NOT–PER–1390–51 (2003).

[39] H.B. Callen and T.A. Welton, *Irreversibility and generalized noise*, *Phys. Rev.* **83** (1951) 34.

[40] P. Amico, L. Bosi, L. Carbone, L. Gammaitoni, F. Marchesoni, et al., *Monolithic fused silica suspension for the Virgo gravitational waves detector*, *Rev. Sci. Instrum.* **73** (2002) 3318.

[41] G. Cagnoli, L. Gammaitoni, J. Kovalik, F. Marchesoni and M. Punturo, *Low-frequency internal friction in clamped-free thin wires*, *Phys. Lett.* **A 255** (1999) 230.

2012 JINST 7 P03012

[42] G.Cagnoli et al., *Mechanical shot noise induced by creep in suspension devices*, *Phys. Lett.* **A 237** (1997) 21.

[43] G. Cagnoli, L. Gammaitoni, J. Kovalik, F. Marchesoni and M. Punturo, *Suspension losses in low-frequency mechanical pendulums*, *Phys. Lett.* **A 213** (1996) 245.

[44] S. Buchman et al., *The Gravity Probe B Relativity Mission*, *Adv. Space Res.* **25** (2000) 1177.

[45] VIRGO collaboration, F. Acernese et al., *First locking of the Virgo central area interferometer with suspension hierarchical control*, *Astropart. Phys.* **20** (2004) 629.

[46] VIRGO collaboration, F. Acernese et al., *A local control system for the test masses of the Virgo gravitational wave detector*, *Astropart. Phys.* **20** (2004) 617.

[47] F. Bellacchia et al., *A VME/VSB board for real-time positioning and image analysis*, LAPP Internal Report (1995).

[48] F. Bellacchia et al., *A VME based CCD imaging system for the Virgo interferometer control*, *Nucl. Instr. Meth. Phys. Res.* **A 413** (1998) 151.

[49] A. Bernardini, E. Majorana, P. Puppo, P. Rapagnani, F. Ricci, et al., *Suspension last stages for the mirrors of the Virgo interferometric gravitational wave antenna*, *Rev. Sci. Instrum.* **70** (1999) 3463.

[50] P. Puppo and P. Rapagnani, *The electromagnetic actuators of the mirror reaction masses*, Virgo Internal Report VIR-NOT-ROM-1390-311 (2005).

[51] VIRGO collaboration, F. Bondu et al., *The VIRGO injection system*, *Class. Quant. Grav.* **19** (2002) 1829.

[52] R. Barillet, A. Brillet, R. Chiche, F. Cleva, L. Latrach and C.N. Man, *An injection-locked Nd:YAG laser for the interferometric detection of gravitational waves*, *Meas. Sci. Technol.* **7** (1996) 162.

[53] http://www.corning.com/specialtymaterials/products_capabilities/ULE.aspx.

[54] F. Bondu, P. Fritschel, C. N. Man, and A. Brillet, *Ultrahigh-spectral purity laser for the VIRGO experiment*, *Opt. Lett.* **21** (1996) 582;
K. Numata, A. Kemery and J. Camp, *Thermal-noise limit in the frequency stabilization of lasers with rigid cavities*, *Phys. Rev. Lett.* **93** (2004) 25.

[55] http://www.schott.com/advanced_optics/english/our_products/zerodur/index.html

[56] F. Beauville et al., *Improvement in the shot noise of a laser interferometer gravitational wave detector by means of an output mode-cleaner*, *Class. Quant. Grav.* **23** (2006) 3235.

[57] F. Beauville et al., *A camera based position control of a suspended optical bench used in a gravitational wave detector*, *Rev. Sci. Instrum.* **74** (2003) 2564.

[58] B. Caron et al., *Photodiodes selection for the VIRGO detector the first step*, *Nucl. Instr. Meth. Phys. Res.* **A 360** (1995) 379.

[59] VIRGO collaboration, R. Flaminio, *The signal detection system for the VIRGO interferometer*, *Nucl. Instrum. Meth.* **A 409** (1998) 477. VIRGO collaboration, R. Flaminio, *Interferometer signal detection system for the VIRGO experiment*, *Class. Quant. Grav.* **19** (2002) 1857.

[60] F. Acernese, P. Amico, M. Al-Shourbagy, S. Aoudia, S. Avino, et al., *The Virgo automatic alignment system*, *Class. Quant. Grav.* **23** (2006) S91.

[61] L. Pinard et al., *95'Optical interference coatings*, *Tucson, Technical Digest* **17** (1995) 200.

[62] J.M. Mackowski et al., *98'Optical interference coatings*, *Tucson, Technical Digest* **9** (1998) 18.

2012 JINST 7 P03012

[63] L. Dognin, *Etude de l'uniformite des multicouches dielectriques des optiques de Virgo par simulation numerique*, *Ph.D. thesis*, Université Claude Bernard Lyon I France, (1997).

[64] M. Bray, *Stitching interferometry: how and why it works*, SPIE **3739** (1999) 259.

[65] M. Bray, *Stitching interferometry and absolute surface shape metrology*, SPIE **4451** (2001) 333.

[66] B. Cimma et al., *Original optical metrologies of large components*, SPIE **5252** (2004) 322.

[67] *RIO2 8062 User Manual*, PowerPC based RISC I/O Boards, DOC 8062/UM, CES (1998).

[68] F. Bellachia, *Virgo Note* VIR-TRE-LAP-5200-101 (1996).

[69] http://www.lynxos.com.

[70] D. Boget et al., *Virgo Note* VIR-SPE-LAP-5200-105 (1998).

[71] D. Boget et al., *Virgo Note* VIR-SPE-LAP-5200-103 (1997).

[72] F. Acernese et al., *The Virgo data acquisition system*, in *XIII IEEE-NPSS Conference on Real Time Systems*, Montréal, Canada, May 18–23 2003.

[73] N. Arnaud, C. Arnault, M. Barsuglia, M. Bizouard, V. Brisson, et al., *The global control of the Virgo experiment*, *Nucl. Instrum. Meth.* **A 550** (2005) 467.

[74] B. Caron, L. Derome, R. Flaminio, X. Grave, F. Marion, et al., *SIESTA, a time domain, general purpose simulation program for the VIRGO experiment*, *Astropart. Phys.* **10** (1999) 369.

[75] F. Acernese, P. Amico, M. Al-Shourbagy, S. Aoudia, S. Avino, et al., *The variable finesse locking technique*, *Class. Quant. Grav.* **23** (2006) S85.

[76] A. Anastasio, F. Barone, A. Eleuteri, F. Garufi, L. Milano, *The Environment Monitoring of the Virgo Antenna for gravitational wave detection*, in *Gravitational Waves*, S. Meshkov ed., proceedings of the 3rd Amaldi Conference, Pasadena, California, (U.S.A), 17–19 July 1999, *AIP Conf. Proc.* **523** (1999) 463.

[77] F. Barone, R. De Rosa, A. Eleuteri, L. Milano and K. Qipiani, *The environmental monitoring system of VIRGO antenna for gravitational waves detection*, *IEEE Trans. Nucl. Sci.* **49** (2002) 405.

[78] VIRGO collaboration, F. Acernese et al., *Properties of seismic noise at the VIRGO site*, *Class. Quant. Grav.* **21** (2004) S433.

[79] F. Acernese et al., *A first study of environmental noise coupling to the Virgo interferometer*, *Class. Quant. Grav.* **22** (2005) S1069.

[80] F. Acernese et al., *Data Acquisition system of the Virgo gravitational waves interferometric detector*, *IEEE Trans. Nucl. Sci.* **55** (2008) 225.

[81] http://www.lapp.in2p3.fr/virgo/FrameL.

[82] B. Mours, Virgo Note VIR-SPE-LAP-5400-102 (1997).

[83] S. Cortese, EGO Computing Center site report, *INFN Computing Workshop 2004,* http://www.infn.it/CCR/workshop/ws2004/doc/II_giornata/cortese_ufficiale.ppt.

[84] A. Bozzi, L. Salconi, VSR2 / S6 Data Transfer report, https://tds.ego-gw.it/itf/tds/file.php?callFile=VIR-0106B-10.ppt.

[85] T. Accadia et al., *Calibration and sensitivity of the Virgo detector during its second science run, Class. Quant. Grav*. **28** (2011) 025005 [arXiv:1009.5190v1].

2012 JINST 7 P03012

[86] L. Blanchet, B.R. Iyer, C.M. Will and A.G. Wiseman, *Gravitational wave forms from inspiralling compact binaries to second postNewtonian order*, *Class. Quant. Grav.* **13** (1996) 575 [gr-qc/9602024];
A. Buonanno, G.B. Cook and F. Pretorius, *Inspiral, merger and ring-down of equal-mass black-hole binaries*, *Phys. Rev.* **D 75** (2007) 124018 [gr-qc/0610122].

[87] F. Marion et al., *Multi-band search of coalescing binaries applied to Virgo CITF data*, in proceedings of *Rencontres deMoriond*, Les Arcs, France, March 2003.

[88] P. Amico et al., *A parallel Beowulf-based system for the detection of gravitational waves in interferometric detectors*, *Comput. Phys. Commun.* **153** (2003) 179.

2012 JINST 7 P03012