ALICE Collaboration

Alessandria, Italy, Facoltà di Scienze dell'Università and INFN:
P. Cortese, G. Dellacasa, L. Ramello and M. Sitta.

Aligarh, India, Physics Department, Aligarh Muslim University:
N. Ahmad, S. Ahmad, T. Ahmad, W. Bari, M. Irfan and M. Zafar.

Athens, Greece, University of Athens, Physics Department:
A. Belogianni, P. Ganoti, A. Petridis, F. Roukoutakis, M. Spyropoulou-Stassinaki and M. Vassiliou.

Bari, Italy, Dipartimento di Fisica dell'Università and Sezione INFN:
G.E. Bruno, M. Caselle, D. Di Bari, D. Elia, R.A. Fini, B. Ghidini, V. Lenti, V. Manzari, E. Nappi, F. Navach, C. Pastore, V. Paticchio, F. Posa, R. Santoro and I. Sgura.

Bari, Italy, Politecnico and Sezione INFN:
F. Corsi, D. De Venuto, U. Fratino and C. Marzocca.

Beijing, China, China Institute of Atomic Energy:
X. Li, Z. Liu, S. Lu, Z. Lu, Q. Meng, B. Sa, J. Yuan, J. Zhou and S. Zhou.

Bergen, Norway, Department of Physics, University of Bergen:
J. Alme, S. Bablok, A. Klovning, J. Nystrand, B. Pommeresche, M. Richter, D. Röhrich, K. Ullaland and H. Yang.

Bergen, Norway, Bergen University College, Faculty of Engineering:
H. Helstrup K. Hetland and K. Røed.

Bhubaneswar, India, Institute of Physics:
R.K. Choudhury, A.K. Dubey, D.P. Mahapatra, D. Mishra, S.C. Phatak and R. Sahoo.

Birmingham, United Kingdom, School of Physics and Space Research, University of Birmingham:
D. Evans, G.T. Jones, P. Jovanović, A. Jusko, J.B. Kinson, R. Lietava and O. Villalobos Baillie.

Bologna, Italy, Dipartimento di Fisica dell'Università and Sezione INFN:
A. Alici, S. Antinori, P. Antonioli, S. Arcelli, G. Bari, M. Basile, L. Bellagamba, D. Boscherini, A. Bruni, G. Bruni, G. Cara Romeo, L. Cifarelli, F. Cindolo, M. Corradi, D. Falchieri, A. Gabrielli, E. Gandolfi, P. Giusti, D. Hatzifotiadou, G. Laurenti, M.L. Luvisetto, A. Margotti, M. Masetti, R. Nania, F. Noferini, F. Palmonari, A. Pesci, A. Polini, G. Sartorelli, E. Scapparone, G. Scioli, G.P. Vacca, G. Valenti, G. Venturi, M.C.S. Williams, C. Zampolli and A. Zichichi.

Bratislava, Slovakia, Comenius University, Faculty of Mathematics, Physics and Informatics:
V. Černý, R. Janik, L. Lúčan, M. Pikna, J. Pišút, N. Pišútová, B. Sitár, P. Strmeň, I. Szarka and M. Zagiba.

Bucharest, Romania, National Institute for Physics and Nuclear Engineering:
C. Aiftimiei, V. Catanescu, M. Duma, C.I. Legrand, D. Moisa, M. Petrovici and G. Stoicea.

Budapest, Hungary, KFKI Research Institute for Particle and Nuclear Physics, Hungarian Academy of Sciences:
E. Dénes, Z. Fodor, T. Kiss, G. Pálla and J. Zimányi.

Cagliari, Italy, Dipartimento di Fisica dell'Università and Sezione INFN:
I. Atanassov S. Basciu, C. Cicalo, A. De Falco, M. Floris, A. Masoni, D. Mura G. Puddu, S. Serci, E. Siddi and G. Usai.

Catania, Italy, Dipartimento di Fisica dell'Università and Sezione INFN:
A. Badalà, R. Barbera, G. Lo Re, A. Palmeri, G.S. Pappalardo, A. Pulvirenti and F. Riggi.

CERN, Switzerland, European Organization for Nuclear Research:
G. Anelli, I. Augustin1), A. Augustinus, J. Baechler, J.A. Belikov2), L. Betev, A. Boccardi, R. Brun, M. Burns, P. Buncic3), I. Cali4), R. Campagnolo, M. Campbell, F. Carena, W. Carena, F. Carminati, N. Carrer, S. Chapeland, C. Cheshkov, P. Chochula, P. Christianssen, P. Christakoglou, A. Colla5), J. Conrad, B.F. Correia Belbute, M. Davenport, G. De Cataldo6), J. de Groot, A. Di Mauro, R. Divià, C. Engster, S. Evrard, C.W. Fabjan, F. Formenti, U. Fuchs, A. Gallas-Torreira, A. Garcia Lopez A. Gheata, M. Gheata7), C. Gonzalez-Gutierrez, R. Grosso5), M. Gruwe8), H.-A. Gustafsson9), H. Hoedelmoser, P. Hristov, M. Ivanov, P. Jacobs10), L. Jirden, A. Junique, S. Kapusta, W. Kickinger, W. Klempt, A. Kluge, T. Kuhr11), L. Leistam, J.P. Lo, M. Lopez Noriega, C. Lourenço, I. Makhlyueva12), J.-C. Marin, P. Martinengo, D. Meunier-Picard, M. Meoni13), M. Morel, A. Morsch, B. Mota, H. Muller, L. Musa, P. Nilsson, D. Nouais14), F. Osmic, D. Perini, A. Peters3), V. Pinto Morais, S. Popescu15), F. Rademakers, J.-P. Revol, P. Riedler, W. Riegler, K. Šafařík, P. Saiz, K. Schossmaier, J. Schukraft, Y. Schutz16), P.A. Silva Loureiro, C. Soos17), G. Stefanini, D. Swoboda, M. Tadel, H. Taureg, M. Tavlet, P. Tissot-Daguette, C. Torcato de Matos, P. Vande Vyvre and J.-P. Vanuxem.

Chandigarh, India, Physics Department, Panjab University:
M.M. Aggarwal, A.K. Bhati, A. Kumar, M. Sharma and G. Sood.

Clermont-Ferrand, France, Laboratoire de Physique Corpusculaire (LPC), IN2P3-CNRS and Université Blaise Pascal:
A. Baldit, V. Barret, N. Bastid, G. Blanchard, J. Castor, P. Crochet, F. Daudon, A. Devaux, P. Dupieux, P. Force, B. Forestier, F. Guerin, R. Guernane, C. Insa, F. Jouve, J. Lecoq, F. Manso, P. Rosnet, L. Royer, P. Saturnini, G. Savinel, G. Stoicea and F. Yermia.

Columbus, U.S.A., Department of Physics, Ohio State University:
T.J. Humanic, I.V. Kotov, M. Lisa, B.S. Nilsen and D. Truesdale.

Columbus, U.S.A., Ohio Supercomputer Centre:
D. Johnson.

Copenhagen, Denmark, Niels Bohr Institute:
I. Bearden, H. Bøggild, C.H. Christensen, J.J. Gaardhøje, K. Gulbrandsen, B.S. Nielsen and G. Renault.

Cracow, Poland, Henryk Niewodniczanski Institute of Nuclear Physics, High Energy Physics Department:
J. Bartke, E. Gładysz-Dziaduś, E. Kornaś, M. Kowalski, A. Rybicki and A. Wróblewski18).

Darmstadt, Germany, Gesellschaft für Schwerionenforschung (GSI):
A. Andronic7), D. Antonczyk, E. Badura, R. Bailhache E. Berdermann, P. Braun-Munzinger, O. Busch, M. Ciobanu7), P. Foka, U. Frankenfeld, C. Garabatos, H. Gutbrod, C. Lippmann, P. Malzacher, A. Marin, D. Miśkowiec, S. Radomski, A. Sandoval, H.R. Schmidt, K. Schwarz, S. Sedykh, R.S. Simon, D. Soyk, H. Stelzer, G. Tziledakis, D. Vranic and J. Wiechula.

Darmstadt, Germany, Institut für Kernphysik, Technische Universität:
U. Bonnes, I. Kraus and H. Oeschler.

Frankfurt, Germany, Institut für Kernphysik, Johann Wolfgang Goethe-Universität:
J. Berger, A. Billmeier, C. Blume, T. Dietel, D. Flierl, M. Gaździcki, Th. Kolleger, S. Lange, C. Loizides, R. Renfordt, R. Stock and H. Ströbele.

Gatchina, Russia, St. Petersburg Nuclear Physics Institute:
Ya. Berdnikov, A. Khanzadeev, N. Miftakhov, V. Nikouline, V. Poliakov, E. Rostchine, V. Samsonov, O. Tarasenkova, V. Tarakanov and M. Zhalov.

Havana, Cuba, Centro de Aplicaciones Technlógicas y Desarrollo Nuclear (CEADEN):
E. Lopez Torres, A. Abrahantes Quintana and R. Diaz Valdes.

Heidelberg, Germany, Kirchhoff Institute for Physics:
V. Angelov, M. Gutfleisch, V. Lindenstruth, R. Panse, C. Reichling, R. Schneider, T. Steinbeck H. Tilsner and A. Wiebalck.

Heidelberg, Germany, Physikalisches Institut, Ruprecht-Karls Universität:
C. Adler, J. Bielčiková, D. Emschermann, P. Glässel, N. Herrmann, Th. Lehmann, W. Ludolphs, T. Mahmoud, J. Milosevic, K. Oyama, V. Petráček, M. Petrovici, I. Rusanov, R. Schicker, H.C. Soltveit, J. Stachel, M. Stockmeier, B. Vulpescu, B. Windelband and S. Yurevich.

Hiroshima, Japan, University of Hiroshima*):
T. Sugitate.

Jaipur, India, Physics Department, University of Rajasthan:
S. Bhardwaj, R. Raniwala and S. Raniwala.

Jammu, India, Physics Department, Jammu University:
S.K. Badyal, A. Bhasin, A. Gupta, V.K. Gupta, S. Mahajan, L.K. Mangotra, B.V.K.S. Potukuchi and S.S. Sambyal.

JINR, Russia, Joint Institute for Nuclear Research:
P.G. Akichine, V.A. Arefiev, Ts. Baatar(19), B.V. Batiounia, V.F. Chepurnov, S.A. Chernenko, V.K. Dodokhov, O.V. Fateev, A.G. Fedunov, M. Haiduc(20), D. Hasegan(20), V.G. Kadychevsky, G. Kharadze(21), B. Khurelbaatar(19), E.K. Koshurnikov, V.L. Lioubochits, V.I. Lobanov, L.V. Malinina, Y.I. Minaev, M. Nioradze(22), P.V. Nomokonov, Y.A. Panebrattsev, V.N. Penev, V.G. Pismennaya, T.A. Pocheptsov, I. Roufanov, G.S. Shabratova, V. Shestakov(23), A.I. Shklovskaya, A.S. Sorin, M.K. Suleimanov, Y. Tevzadze(22), R. Togoo(19), A.S. Vodopianov, V.I. Yurevich, Y.V. Zanevsky, S.A. Zaporojets and A.I. Zinchenko.

Jyväskylä, Finland, Department of Physics, University of Jyväskylä and Helsinki Institute of Physics:
J. Äysto, M. Bondila, V. Lyapin, M. Oinonen, T. Malkiewicz, V. Ruuskanen, H. Seppänen, W. Trzaska and S. Yamaletkinov.

Kangnung, Republic of Korea, Kangnung National University:
H.T. Jung, W. Jung, D.-W. Kim, H.N. Kim, J.S. Kim, K.S. Lee and S.-C. Lee

Karlsruhe, Germany, Institut für Prozessdatenverarbeitung und Elektronik (IPE)*):
T. Blank and H. Gemmeke.

Kharkov, Ukraine, National Scientific Centre, Kharkov Institute of Physics and Technology:
G.L. Bochek, A.N. Dovbnya, V.I. Kulibaba, N.I. Maslov, S.V. Naumov, V.D. Ovchinnik, S.M. Potin and A.F. Starodubtsev.

Kharkov, Ukraine, Scientific and Technological Research Institute of Instrument Engineering:
V.N. Borshchov, O. Chykalov, L. Kaurova, S.K. Kiprich, L. Klymova, O.M. Listratenko, N. Mykhaylova, M. Protsenko, O. Reznik and V.E. Starkov.

Kiev, Ukraine, Department of High Energy Density Physics, Bogolyubov Institute for Theoretical Physics, National Academy of Sciences of Ukraine:
O. Borysov, I. Kadenko, Y. Martynov, S. Molodtsov, O. Sokolov, Y. Sinyukov and G. Zinovjev.

Kolkata, India, Saha Institute of Nuclear Physics:
P. Bhattacharya, S. Bose, S. Chattopadhyay, N. Majumdar, S. Mukhopadhyay, A. Sanyal, S. Sarkar, P. Sen, S.K. Sen, B.C. Sinha and T. Sinha.

Kolkata, India, Variable Energy Cyclotron Centre:
Z. Ahammed, P. Bhaskar, S. Chattopadhyay, D. Das, S. Das, M.R. Dutta Majumdar, M.S. Ganti, P. Ghosh, B. Mohanty, T.K. Nayak, P.K. Netrakanti, S. Pal, R.N. Singaraju, V. Singhal, B. Sinha, M.D. Trivedi and Y.P. Viyogi.

Köln, Germany, University of Applied Sciences Cologne, Communications Engineering*):
G. Hartung and T. Krawutschke.

Košice, Slovakia, Institute of Experimental Physics, Slovak Academy of Sciences and Faculty of Science, P.J. Šafárik University:
J. Bán, M. Bombara, A. Dirner, S. Fedor, M. Hnatič, I. Králik, A. Kravčáková, F. Kriváň, M. Krivda, G. Martinská, B. Pastirčák, L. Šándor, J. Urbán and J. Vrláková.

Legnaro, Italy, Laboratori Nazionali di Legnaro:
M. Cinausero, E. Fioretto, G. Prete, R.A. Ricci and L. Vannucci.

Lisbon, Portugal, Departamento de Fìsica, Instituto Superior Tècnico:
P. Branco, R. Carvalho, J. Seixas and R. Vilela Mendes.

Lund, Sweden, Division of Experimental High Energy Physics, University of Lund:
A. Oskarsson, L. Osterman, I. Otterlund and E.A. Stenlund.

Lyon, France, Institut de Physique Nucléaire de Lyon (IPNL), IN2P3-CNRS and Université Claude Bernard Lyon-I:
B. Cheynis, L. Ducroux, J.Y. Grossiord, A. Guichard, P. Pillot, B. Rapp and R. Tieulent.

Mexico City and Merida, Mexico, Centro de Investigacion y de Estudios Avanzados del IPN, Universidad Nacional Autonoma de Mexico, Instituto de Ciencias Nucleares, Instituto de Fisica:
J.R. Alfaro Molina, A. Ayala, E. Belmont Moreno, J.G. Contreras, E. Cuautle, J.C. D'Olivo, I. Dominguez, A. Flores, G. Herrera Corral, I. Leon Monzon, M. Linares, J. Martinez Castro, M.I. Martinez, A. Martinez Davalos, A. Menchaca-Rocha, L.M. Montano Zetina, L. Nellen, G. Paic(24), J. del Pino, P. Reyes, A. Sandoval, J. Solano, M.A. Vargas and A. Zepeda.

Moscow, Russia, Institute for Nuclear Research, Academy of Science:
V.A. Feshcenko, M.B. Golubeva, V.G. Gorlychev, F.F. Guber, O.V. Karavichev, T.L. Karavicheva, E.V. Karpechev, A.B. Kurepin, A.I. Maevskaya, V.V. Marin, I.A. Pshenichnov, V.I. Razin, A.I. Rechetin, K.A. Shileev and N.S. Topil'skaia.

Moscow, Russia, Institute for Theoretical and Experimental Physics:
A.N. Akindinov, V. Golovine, A.B. Kaidalov, M.M. Kats, I.T. Kiselev, S.M. Kisselev, E. Lioublev, M. Martemianov, A.N. Martemiyanov, P.A. Polozov, V.S. Serov, A.V. Smirnitski, M.M. Tchoumakov, I.A. Vetlitski, K.G. Volochine, L.S. Vorobiev and B.V. Zagreev.

Moscow, Russia, Russian Research Center Kurchatov Institute :
D. Aleksandrov, V. Antonenko, S. Beliaev, S. Fokine, M. Ippolitov, K. Karadjev, V. Lebedev, V.I. Manko, T. Moukhanova, A. Nianine, S. Nikolaev, S. Nikouline, O. Patarakine, D. Peressounko, I. Sibiriak, A. Tsvetkov, A. Vasiliev, A. Vinogradov, M. Volkov and I. Yushmanov.

Moscow, Russia, Moscow Engineering Physics Institute:
V.A. Grigoriev, V.A. Kaplin and V.A. Loginov.

Mumbai, India, Indian Institute of Technology,*):
B.K. Nandi and R. Varma.

Mumbai, India Bhabha Atomic Research Centre (BARC), *):
V. Chandratre and V. Kataria.

Münster, Germany, Institut für Kernphysik, Westfälische Wilhelms Universität:
C. Baumann, D. Bucher, R. Glasow, H. Gottschlag, J.F. Gros̆e-Oetringhaus, N. Heine, C. Klein-Bösing, K. Reygers, R. Santo, W. Verhoeven, J. Wessels and A. Wilk.

Nantes, France, Laboratoire de Physique Subatomique et des Technologies Associées (SUBATECH), Ecole des Mines de Nantes, IN2P3-CNRS and Université de Nantes:
L. Aphecetche, J.M. Barbet, G. Conesa-Balbastre, Z. Conesa-del-Valle, J.P. Cussonneau, H. Delagrange, M. Dialinas, C. Finck, B. Erazmus, M. Germain, S. Kabana, F. Lefèvre, L. Luquin, L. Martin, G. Martinez, C. Roy and A. Tournaire.

Novosibirsk, Russia, Budker Institute for Nuclear Physics:
A.R. Frolov and I.N. Pestov.

Oak Ridge, U.S.A., Oak Ridge National Laboratory:
T. Awes.

Omaha, U.S.A., Creighton University:
M. Cherney.

Orsay, France, Institut de Physique Nucléaire (IPNO), IN2P3-CNRS and Université de Paris-Sud:
L. Bimbot, V. Chambert, A. Charpy, M.P. Comets, P. Courtat, S. Drouet, P. Edelbruck, B. Espagnon, I. Hřivnáčová, R. Kunne, Y. Le Bornec, M. Mac Cormick, J. Peyré, J. Pouthas, S. Rousseau, C. Suire and N. Willis.

Oslo, Norway, Department of Physics, University of Oslo:
L. Bravina, G. Løvhøiden, B. Skaali, T.S. Tveter and T. Vik.

Padua, Italy, Dipartimento di Fisica dell'Università and Sezione INFN:
F. Antinori, A. Dainese, R. Dima, D. Fabris, J. Faivre, M. Lunardon, M. Morando, S. Moretto, A. Pepato, E. Quercigh, F. Scarlassara, G. Segato, R. Turrisi and G. Viesti.

Pohang, Republic of Korea, Pohang Accelerator Laboratory:
J. Choi.

Prague, Czech Republic, Institute of Physics, Academy of Science:
A. Beitlerova, J. Mareš, K. Polák and P. Závada.

Prague, Czech Republic, Czech Technical University Faculty of Nuclear Sciences and Physical Engineering:
V. Petráček, M. Pachr and L. Škoda.

Protvino, Russia, Institute for High Energy Physics:
M.Yu. Bogolyubsky, G.V. Khaoustov, I.V. Kharlov, N.G. Minaev, V.S. Petrov, B.V. Polichtchouk, S.A. Sadovsky, V.A. Senko, A.S. Soloviev, P.V. Stolpovsky and V.A. Victorov.

Puebla, Mexico, Benemerita Universidad Autonoma de Puebla:
A. Fernandez Tellez, E. Gamez Flores, R. Lopez, G. Tejeda and S. Vergara

Řež u Prahy, Czech Republic, Academy of Sciences of Czech Republic, Nuclear Physics Institute:
D. Adamová, S. Kouchpil, V. Kouchpil, A. Kugler, M. Šumbera and V. Wagner.

Rome, Italy, Dipartimento di Fisica dell'Università 'La Sapienza' and Sezione INFN:
S. Di Liberto, M.A. Mazzoni, F. Meddi and G.M. Urciuoli.

Rondebosch, South Africa, University of Cape Town:
J. Cleymans, G. de Vaux, R W. Fearick, A. Szostak and Z.Z. Vilakazi.

Saclay, France, Centre d'Études Nucléaires, DAPNIA:
M. Anfreville, A. Baldisseri, B. Becker, H. Borel, J.-L. Charvet, M. Combet, J. Gosset, P. Hardy, S. Herlant, F. Orsini, Y. Pénichot, H. Pereira, S. Salasca, F.M. Staley and M. Usseglio.

Salerno, Italy, Dipartimento di Fisica 'E.R.Caianiello' dell'Università and Sezione INFN:
A. De Caro, C. Guarnaccia, D. De Gruttola, S. De Pasquale, A. Di Bartolomeo, M. Fusco Girard, G. Grella, M. Guida, G. Romano, S. Sellitto, R. Silvestri and T. Virgili.

Sarov, Russia, Russian Federal Nuclear Center (VNIIEF):
V. Basmanov, D. Budnikov, V. Demanov, V. Ianowski, R. Ilkaev, L. Ilkaeva, A. Ivanov, A. Khlebnikov, A. Kouryakin, S. Nazarenko, V. Pavlov, S. Philchagin, V. Punin, S. Poutevskoi, I. Selin, I. Vinogradov, S. Zhelezov and A. Zhitnik.

Split, Croatia, Technical University of Split, Faculty of Electrical Engineering, Mechanical Engineering and Naval Architecture (FESB):
S. Gotovac, E. Mudnic and L. Vidak.

St. Petersburg, Russia, V. Fock Institute for Physics of St. Petersburg State University :
M.A. Braun, G.A. Feofilov, S.N. Igolkin, A.A. Kolojvari, V.P. Kondratiev, P.A. Otyugova, O.I. Stolyarov, T.A. Toulina, F.A. Tsimbal, F.F. Valiev, V.V. Vechernin and L.I. Vinogradov.

Strasbourg, France, Institut de Recherches Subatomiques (IReS), IN2P3-CNRS and Université Louis Pasteur:
J. Baudot, D. Bonnet, J.P. Coffin, B. Hippolyte, C. Kuhn, J.R. Lutz and R. Vernet.

Tokyo, Japan, University of Tokyo,*):
H. Hamagaki and K. Ozawa.

Trieste, Italy, Dipartimento di Fisica dell'Università and Sezione INFN:
V. Bonvicini, O. Borysov, L. Bosisio, M. Bregant, P. Camerini, G. Contin, F. Faleschini, E. Fragiacomo, N. Grion, G. Margagliotti, S. Piano, I. Rachevskaya, A. Rachevski, R. Rui and A. Vacchi.

Turin, Italy, Dipartimenti di Fisica dell'Università and Sezione INFN:
B. Alessandro, R. Arnaldi, G. Batigne, S. Beolé, E. Bruna, P. Cerello, E. Chiavassa, S. Coli, E. Crescio, N. De Marco, A. Ferretti, M. Gagliardi, M. Gallio, L. Gaudichet, R. Gemme, G. Giraudo, P. Giubellino, M. Idzik, S. Martoiu, A. Marzari Chiesa, M. Masera, G. Mazza, P. Mereu, M. Monteno, A. Musso, C. Oppedisano, A. Piccotti, F. Poggio, F. Prino, L. Riccati, A. Rivetti, E. Scomparin, F. Tosello, L. Toscano, G. Travaglia, E. Vercellin, A. Werbrouck and R. Wheadon.

Tsukuba, Japan, University of Tsukuba,*):
Y. Miake and S. Esumi.

Utrecht, The Netherlands, Subatomic Physics Department, Utrecht University and National Institute for Nuclear and High Energy Physics (NIKHEF):
M. Botje, J.J.F. Buskop, A.P. De Haas, R. Kamermans, P.G. Kuijer, G. Nooren, C.J. Oskamp, Th. Peitzmann, E. Simili, R. Snellings, A.N. Sokolov, A. Van Den Brink and N. Van Eijndhoven.

Wako-shi, Japan, Institute of Research, RIKEN*):
H. Enyo, K. Fujiwara, H. Kano and H. Onishi.

Warsaw, Poland, Soltan Institute for Nuclear Studies:
A. Deloff, T. Dobrowolski, K. Karpio, M. Kozlowski, H. Malinowski, K. Redlich25), T. Siemiarczuk, G. Stefanek26), L. Tykarski and G. Wilk.

Warsaw, Poland, University of Technology, Institute of Physics:
Z. Chajecki, H. Gos, M. Janik, M. Jedynak, A. Kisiel, T.J. Pawlak, W.S. Peryt, J. Pluta, P. Skowronski, M. Slodkowski, P. Szuba and T. Traczyk.

Worms, Germany, University of Applied Sciences Worms, ZTT*):
E.S. Conner and R. Keidel.

Wuhan, China, Institute of Particle Physics, Huazhong Normal University:
X. Cai, X.R. Wang, T. Wu, H.Y. Yang, Z.B. Yin and D.C. Zhou.

Wuhan, China, Huazhong Univ. of Science and Techn. Electronic and Inform. Engineering:
X. Cao and Q. Li.

Yerevan, Armenia, Yerevan Physics Institute:
M. Atayan, A. Grigoryan, S. Grigoryan, H. Gulkanyan, A. Harutyunyan, A. Hayrapetyan, V. Kakoyan, M. Poghosyan and G. Sargsyan.

Zagreb, Croatia, Ruder Bošković Institute:
T. Anticic, K. Kadija and T. Susa.

---

[*] Associate member.
[1] Now at Darmstadt, Germany, Gesellschaft für Schwerionenforschung (GSI)
[2] On leave from JINR, Dubna, Russia.
[3] supported by EGEE
[4] On leave from Università del Piemonte Orientale, Alessandria, Italy.
[5] On leave from INFN, Italy
[6] also at INFN, Bari, Italy
[7] On leave from National Institute for Physics and Nuclear Engineering, Bucharest, Romania.
[8] Also at GSI, Darmstadt, Germany
[9] On leave from Division of Experimental High Energy Physics, University of Lund, Lund, Sweden
[10] On leave from Lawrence Berkeley National Laboratory
[11] Now at Institut fuer Theoretische Teilchenphysik, Universitaet Karlsruhe
[12] Also at ITEP, Moscow, Russia
[13] Now at INFN-Firenze
[14] Also at Also at INFN, Sezione di Torino, Italy
[15] Also at IFIN-HH, HEP Dept., Bucharest, Romania
[16] On leave from Laboratoire de Physique Subatomique et des Technologies Associées (SUBATECH), Ecole des Mines de Nantes, IN2P3-CNRS and Université de Nantes, Nantes, France.
[17] On leave from Budapest University, Budapest, Hungary.
[18] Cracow Technical University, Cracow, Poland.
[19] Institute of Physics and Technology, Mongolian Academy of Sciences, Ulaanbaatar, Mongolia.
[20] Institute of Space Sciences, Bucharest, Romania.
[21] Institute of Physics, Georgian Academy of Sciences, Tbilisi, Georgia.
[22] High Energy Physics Institute, Tbilisi State University, Tbilisi, Georgia.
[23] Research Centre for Applied Nuclear Physics (RCANP), Dubna, Russia.
[24] Department of Physics, Ohio State University, Columbus, U.S.A.
[25] Physics Faculty, University of Bielefeld, Belfield, Germany and University of Wroclaw, Wroclaw, Poland.
[26] Institute of Physics, Pedagogical University, Kielce, Poland.

Acknowledgements

The Collaboration wishes to thank all the administrative and technical staff involved during the preparation of the Computing TDR and in particular: the ALICE secretariat, M. Connor, U. Genoud, and C. Hervet; the CERN Desktop Publishing Service, in particular S. Leech O'Neale and C. Vanoli; and the CERN printshop.

Contents

7 Computing model and capacity requirements

7.1 Introduction

This chapter describes the computing model for the processing of the data produced by the ALICE detector in pp and A–A collisions every year. The computing model of the LHC experiements has been already reviewed in 2001 during the so-called LHC Computing Review (LHCCR) [1]. The model presented in this document however contains several changes with respect to the 2001 version, and it has been reviewed by a subsequent LHCC review that focussed mainly on the requested resources [2]. Compared to the 2001 version, the model has been further developed and refined thanks to the experience gained during the Physics Data Challenges. The amount of permanent storage has increased because a duplication of raw data at external Tier 1s is now foreseen in the model. Disk storage has substantially increased: the new estimate is deduced from the disk usage made during the Physics Data Challenges. This document contains also an evaluation of the contributions of Tier 2 centres, which was not present in the ALICE LHCC estimates.

7.2 Input parameters

The input parameters for the present computing model are derived from information contained in the ALICE Trigger, DAQ, HLT and Control System TDR [3] and the ALICE Physics Performance Report Vol. 1 [4]. All input parameters are to be considered as our best estimates at the moment. They are based on the nominal figures of a standard data-taking year. The resources requirements during the commissioning of the machine and the two first years of running have been estimated as a percentage of a standard data-taking year, taking into account the information available at this moment about the LHC machine schedule.

Independently from the LHC pp luminosity, the beam parameters will be adjusted to achieve a luminosity of at most 3×10^{30} cm^{-2}s^{-1} leading to an interaction of 200 kHz and an average pileup of about 36 events during the TPC drift time (88 μs). The parameters adopted for Pb–Pb collisions are listed in Table 7.1.

Table 7.1: Parameters adopted for Pb–Pb collisions.

Centre-of-mass energy	5.5A TeV
Luminosity	5×10^{25} cm^{-2}s^{-1} in 2008
	5×10^{26} cm^{-2}s^{-1} (average luminosity) from 2009 onward
Total reaction cross-section	8 b
Nominal collision rate	4×10^{3} Hz Pb–Pb collisions at average luminosity

The trigger rate and RAW data bandwidth are independent of the luminosity as trigger thresholds and conditions will be adjusted to saturate the bandwidth. During pp collisions, more than one bunch crossing will happen during the gating of the TPC on account of its long drift time, and therefore more than one collision may be recorded (pileup). The amount of this pileup depends on luminosity, and therefore the event size will increase with the luminosity. Studies are under way to remove these pileup events on line with the HLT system. However, given the criticality of this operation, we do not foresee doing it in the first years of pp running.

During the first year of operation, the following running scenario has been assumed. The first pp run will take place over a reduced time, from the beginning of July 2007 until the end of September 2007. Assuming the maximum luminosity allowed for ALICE, this first run will deliver 40% of the pp data during a standard data-taking year. This run will be followed by a short Pb–Pb pilot run with reduced luminosity which will deliver at most 20% of the Pb–Pb data we will collect in one standard data-taking year. The first Pb–Pb run with a still reduced luminosity will occur at the end of 2008, the second year of LHC operation. As we have explained this will not imply a reduction in the data rate, as we will adjust the trigger accordingly. These assumptions have been taken into account to establish the rampup of resources deployment required by ALICE.

Below we shall discuss the following types of data:

- **RAW**: raw data as recorded by the DAQ (real) or simulated (Monte Carlo, MC);
- **ESD**: Event Summary Data as produced by the reconstruction program;
- **AOD**: Physics Analysis Object Data derived from the ESD and similar to today's n-tuples;
- **TAG**: Event tags for event selection.

As far as the basic RAW data parameters are concerned, for both pp and Pb–Pb an average acquisition rate is considered. During a pp run, the rate can go up to 500 Hz, however, this is in special cases and for a short period of time only. Since for Pb–Pb collisions the event size changes significantly with the centrality of the collision, a normalized rate, taking an average event size of 12.5 MB, is considered. This event size has been calculated assuming a charged-particle density of $\mathrm{d}N/\mathrm{d}y = 4000$. The actual trigger rate will be the sum of several event types with different sizes and rates.

The data taking parameters are summarised in Table 7.2.

Table 7.2: Data-taking parameters.

	pp	Pb–Pb
Event recording rate (Hz)	100	100
Event recording bandwidth (MB/s)	100	1250
Running time per year (Ms)	10	1
Events per year	10^9	10^8

7.3 The computing model

The computing model is subject to change with time. During the first years of running, even if luminosities might be below the nominal luminosity, only slightly less data will be recorded than during runs with nominal luminosities. The reason is that selective triggers will become operational only after a necessary training period needed to understand the various detection systems. Data from this early period will be processed offline rather differently than data from later runs. Understanding the detectors and training of the alignment, calibration, reconstruction and analysis algorithms will required many processing passes by many physicists on a subset of raw data. This phase must be limited in time and followed by an early fast processing of the total set of available data to guarantee early access to physics results. This fast processing must, however, be complete to preserve the richness of the data and the discovery potential in the early run. Therefore, adequate computing must already be available when the LHC provides the first pp collisions.

In the following we discuss an offline processing scenario, from which the required computing resources are estimated, and which we foresee to apply during standard data-taking years at nominal luminosities. We assume that a normal run period is split into:

- seven months (effective 10^7 s) of pp collisions,
- one month (effective 10^6 s) of heavy-ion collisions, and
- four months of winter shutdown.

The overall organized processing (calibration, alignment, reconstruction and scheduled analysis) is scheduled and prioritized in the ALICE Physics Working Groups (PWG) and by the ALICE Physics Board. The end-user analysis (sometimes called ‘chaotic analysis’) is performed in the context of the PWGs and, in case of lack of resources, it is prioritized by the Physics Board.

Depending on the middleware that will be in use at the time of LHC start-up, a more or less hierarchical organization of computing resources (Tier 0, Tier 1, Tier 2) will be in order. Although ALICE believes that most likely the democratic cloud model rather than the strict hierarchy of the MONARC model will prevail, the estimate of the required resources uses the concepts and names introduced by MONARC (see also Section **??**).

- Tier 0 provides permanent storage for the raw data, distributes raw data to Tier 1 and performs the calibration and alignment task and the first reconstruction pass. Our model foresees that one Tier 1 and one Tier 2 be co-located with the Tier 0 at CERN.
- Tier 1s outside CERN provide permanent storage of a copy of the raw data. All Tier 1s perform the subsequent reconstruction passes, the scheduled analysis tasks, and the reconstruction of Pb–Pb MC data; they have the responsibility for the long term storage of data processed at Tier 1s and Tier 2s.
- Tier 2s generate and reconstruct the simulated Monte Carlo data and perform the chaotic analysis.
- Tier 0, Tier 1s and Tier 2s provide short-term storage with fast access to multiple copies of active data, raw (only a fraction) and processed.

Although the scenario may change to adjust to the available computing resources and Grid middleware functionality at a given time, or to provide rapid feedback to physics results obtained during the initial reconstruction and analysis, it is unlikely that the estimated needs of computing resources will vary considerably. Changes of the order of up to 50% in processing, short-term storage, or long-term storage must however be anticipated. Unexpected features in the operation of the Grid resources, which might result in efficiencies lower than anticipated, and which could not be predicted during the Data Challenges, could also require more or less important changes in the needed computing resources.

To estimate the computing resources required by the present model, the efficiency factors for the usage of processors (CPU), short-term storage (disk), and long-term mass storage (MS) listed in Table 7.3 have been adopted. These are factors agreed by the LCG project and used by all the LHC experiments.

7.4 CPU requirements

7.4.1 Parameter values

The ALICE Data Acquisition system will record pp and heavy-ion data at a rate of 100 Hz during the previously discussed effective time of 10^7 s and 10^6 s, respectively (see Table 7.2). The raw data size will vary with the trigger selection. The product trigger-rate times raw-data-size is limited by the maximum bandwidth available to transfer data to the mass storage system, 1.25 GB/s.

Table 7.3: Efficiency factors for the use of processors (CPU), short-term storage (disk), and long-term mass storage (MS).

Efficiency factors (%)	
Efficiency for scheduled CPU	85
Efficiency for chaotic CPU	60
Disk utilization efficiency	70
MS utilization efficiency	100

The processing power required for the reconstruction of data, raw as well as Monte Carlo, has been estimated from the performance of the reconstruction algorithms presently in use and tested on Monte Carlo data. We have taken as a value 80% of the present reconstruction times, which takes into account the optimization of the code and its increase in complexity within a reasonable safety factor. For Pb–Pb collisions an average charged-particle multiplicity of 4000 has been adopted. This is a conservative estimation based on the recent results from RHIC and the theoretical extrapolations at LHC energy. The real value might, however, be smaller as suggested by the RHIC data. The reconstruction algorithms are not yet entirely optimized, and they are not complete as they do not yet include calibration or alignment procedures. The figures quoted for reconstruction try to take these factors into account, considering the evolution of the code performance during the past years.

During the first years there will be a full first reconstruction pass followed by a 'second pass' which will in reality be composed of several short runs to tune the reconstruction programs based on the results of the first full pass. It is expected that the resources needed by this activity will correspond to a full reconstruction pass. Then we will run a second full reconstruction pass on all data of that year's run with final condition information. As time passes, we expect to improve our capability to derive good condition information soon after the first reconstruction pass. However, we also expect the need to reconstruct and analyse data coming from previous runs. So the estimation of three reconstruction passes is reasonable over the foreseeable lifetime of the experiment, if compound with the foreseen upgrade of the computing equipment (see Section 7.7).

The computing power required for analysis has been estimated based on a limited variety of Grid-enabled analysis use-cases performed during the Physics Data Challenges. Since the analysis algorithms will vary in complexity depending on the physics channel under study, the computing power for analysis will have to be revised once a larger set of analysis classes has been exercised on the Grid. A variation of a factor up to two must be anticipated.

The computing power required for the generation of Monte Carlo data is better under control. It will be subject to changes only if the real charged-particle multiplicity differs substantially from the predicted by the HIJING Monte Carlo generator. Table 7.4 lists the values adopted for the various parameters entering our estimation. The CPU power quoted for simulation includes event generation, tracking of particles through the detectors (an average value for the GEANT 3 and FLUKA transport codes has been adopted), and digitization of the generated signals.

Chaotic analysis tasks are well-focused analyses performed by single users on a subset of the data, possibly residing entirely at given Tier 2 centres. Scheduled analyses tasks gather many analysis from users and thus require much more CPU power than single-analysis tasks. They are performed on the entire set of reconstructed data, once per reconstruction pass. To optimize data transfer, scheduled analysis is best performed at Tier 1 centres.

Table 7.4: Computing power required for the processing (simulation includes event generation, tracking and digitization) of the ALICE real and Monte Carlo data.

Processing power parameters			
		pp	HI
Reconstruction	KSI2k×s/event	5.9	740.0
Chaotic analysis	KSI2k×s/event	0.6	8.3
Scheduled analysis	KSI2k×s/event	16.0	240.0
Simulation	KSI2k×s/event	39.0	17000.0
Reconstruction passes	–	3	3
Chaotic analysis passes	–	20	20
Scheduled analysis passes	–	3	3

7.4.2 Raw-data processing strategy

The scheduled processing of raw data consists of the reconstruction of the raw data and the creation of ESD objects.

Although the processing strategies vary between pp and heavy-ion runs, they have in common that the first reconstruction pass must be fast and must start immediately after (for heavy-ion) or during (for pp) data taking to allow for rapid discoveries and to establish quickly the overall properties of the collisions. On the other hand, the first heavy-ion reconstruction pass must be finished well before the start of the next heavy-ion run (at least six months before) to be able to define from the data analysis the running conditions. For heavy-ion runs, this first reconstruction is preceded by a computing-intensive calibration and alignment task not exceeding one month in duration and taking place during heavy-ion data acquisition. The additional reconstruction passes will consist of a full-fledged reconstruction including fine tuning of all the parameters of the algorithms.

Calibration and alignment tasks are performed quasi online on the events stored on a disk buffer at Tier 0. The size of this buffer is the equivalent of two days of heavy-ion data taking, i.e. 200 TB.

7.4.2.1 pp data processing

The data collected during pp runs, being less demanding in terms of computing resources than heavy-ion data for their reconstruction, will be processed online or quasi online during data taking (seven months according to our standard data-taking year) at the CERN Tier 0. The data temporarily stored on the Tier 0 disk buffer (one day of Pb–Pb running is the equivalent of 10 days pp data taking) are processed, successively calibrated, and reconstructed. Reconstructed data will be kept for long-term storage locally at Tier 0 and distributed to the Tier 1s. The second and the third reconstruction passes will be distributed in Tier 1s and spread over a period of six months.

7.4.2.2 Heavy-ion data processing

Because of the much larger data recording rate in heavy-ion mode, online reconstruction of the entire set of data would require unaffordable computing resources. The first reconstruction pass will therefore last four months and can start immediately after the data taking ends, or even during the run depending on the performance of the calibration and alignment task. It is crucial that this first reconstruction ends well before the next heavy-ion run starts in order to allow for enough time to analyse the data and elaborate the new running conditions. We estimate that the time required for this is of the order of six months. The second and third reconstruction passes will be distributed in Tier 1s and can be done over a period of six months. Obviously, in such a scenario, pp and A–A reconstruction passes will be concurrent. Two different reconstruction tasks, one for pp and one for A–A, will be permanently active in Tier 1s, and at Tier 0 either the first pp or the first A–A reconstruction will be active (see the sketch of a possible processing scenario in Fig. 7.1).

The resources required to perform the ALICE pp and heavy-ion data reconstruction within such a scenario are listed in Table 7.5.

Table 7.5: Annual processing resources required to reconstruct and analyse the real pp and A–A data and to process (generation, reconstruction, and analysis) Monte Carlo data.

	CPU (MSI2K)		
	Integrated	Average	Max.
MC (simulation+reconstruction)	150	12.0	12.0
Real reconstruction	94	7.8	13.0
Scheduled analysis	96	8.0	8.4
Chaotic analysis	16	1.3	1.3
Alignment & calibration	3	0.3	3.0

7.4.3 Monte Carlo data simulation

The production scheme of Monte Carlo differs for pp and heavy-ion simulations. In the case of pp simulations, Monte Carlo data will be generated in an amount similar to that of collected real data (10^9 events per year). To avoid requesting a prohibitive amount of resources and also to enrich generated events with physics signal of interest, we have adopted a merging technique for heavy-ion simulation. Full-fledged, heavy-ion events are generated in a limited amount (10^7 events per year) in sets of different impact parameters (minimum-bias, peripheral, and central); these events are called the ‘underlying events’. The different signal events are then generated, merged into one underlying event, at the digits level, and the merged event is reconstructed in a single task. In this way, the underlying events can be reused several times. From the data challenges experience we have fixed to 10 the number of times an underlying event can be reused and preserve realistic event-to-event fluctuations.

The Monte Carlo pp and heavy-ion events are typically generated and reconstructed at the Tier 2 centres.

The annual resources required to process the Monte Carlo data (generation, reconstruction) are listed in Table 7.5.

7.4.4 Data analysis

Based on the experience of past experiments, the data analysis will consume a large fraction of the total amount of resources. At variance with the reconstruction and the simulation tasks, not all the analysis tasks can be easily scheduled as they will involve potentially all the physicists of the collaboration applying customized algorithms many times on subsets of the data of various sizes. This way of performing analysis is called chaotic analysis. However, it is foreseen to perform scheduled analysis during the same production tasks as the reconstruction passes and therefore over the entire set of data. This scheduled analysis will group in each pass all the official algorithms approved by the Physics Board. The time needed to analyse reconstructed events depends very much on the complexity of the analysis algorithm and on the overhead introduced by the Grid processing. The latter can only be predicted once the analysis on the Grid has been exercised with the final middleware and at the final scale. The processing power we have considered is an average value which can vary largely from one analysis to the other. To estimate

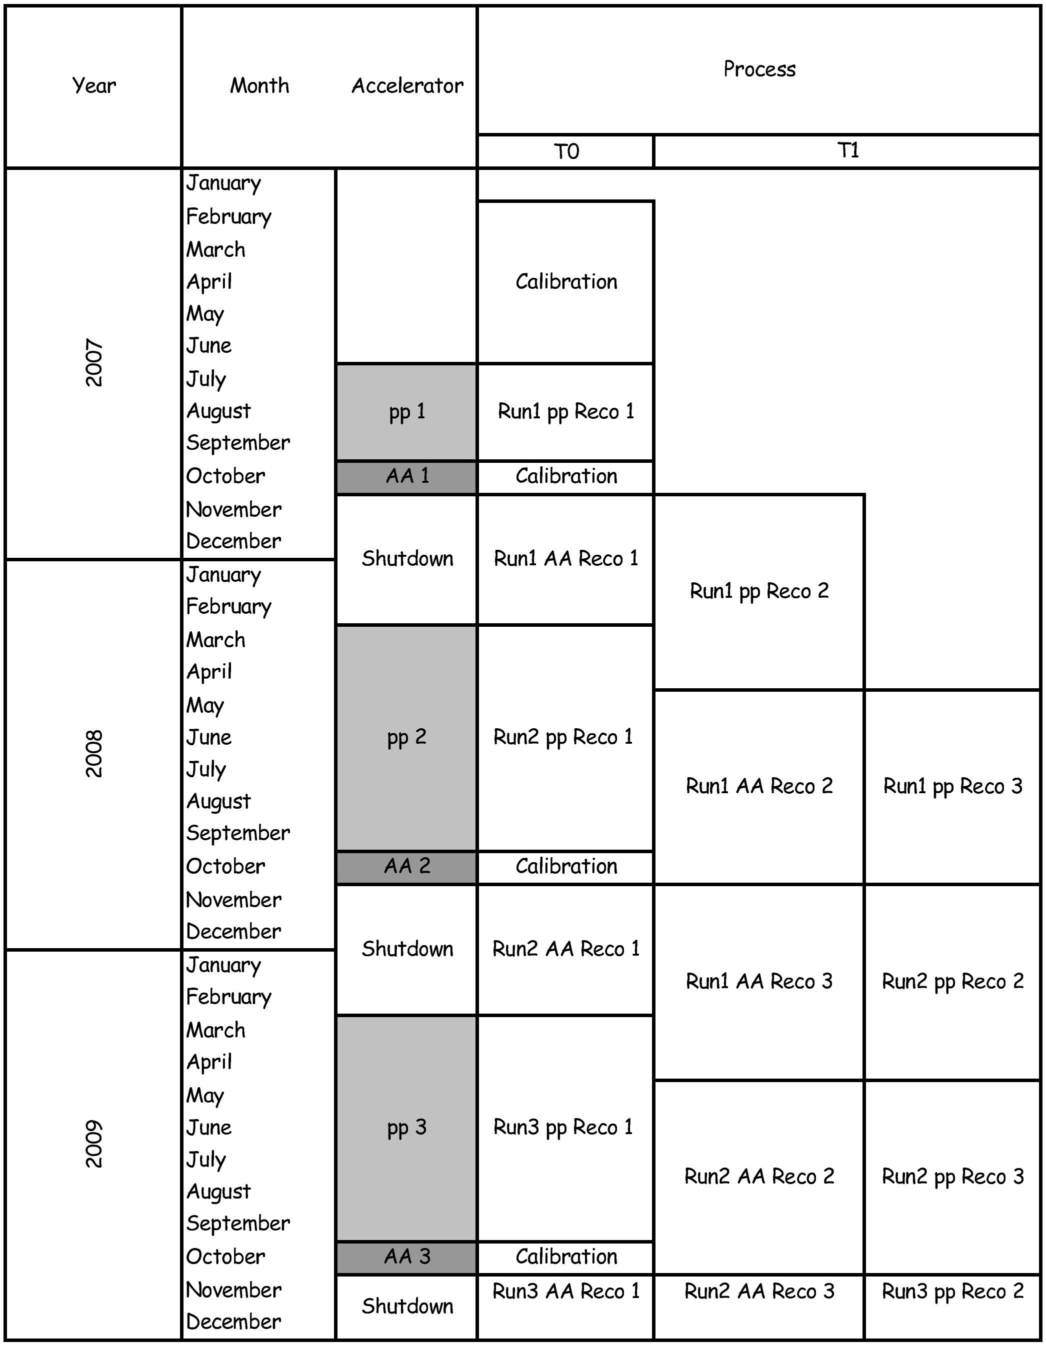

Year	Month	Accelerator	Process		
			T0	T1	
2007	January				
	February		Calibration		
	March				
	April				
	May				
	June				
	July	pp 1	Run1 pp Reco 1		
	August				
	September				
	October	AA 1	Calibration		
	November	Shutdown	Run1 AA Reco 1	Run1 pp Reco 2	
	December				
2008	January				
	February				
	March	pp 2	Run2 pp Reco 1		
	April				
	May			Run1 AA Reco 2	Run1 pp Reco 3
	June				
	July				
	August				
	September				
	October	AA 2	Calibration		
	November	Shutdown	Run2 AA Reco 1	Run1 AA Reco 3	Run2 pp Reco 2
	December				
2009	January				
	February				
	March	pp 3	Run3 pp Reco 1		
	April				
	May			Run2 AA Reco 2	Run2 pp Reco 3
	June				
	July				
	August				
	September				
	October	AA 3	Calibration		
	November	Shutdown	Run3 AA Reco 1	Run2 AA Reco 3	Run3 pp Reco 2
	December				

Figure 7.1: Processing scenario used to estimate the resources required at Tier 0, Tier 1s and Tier 2s.

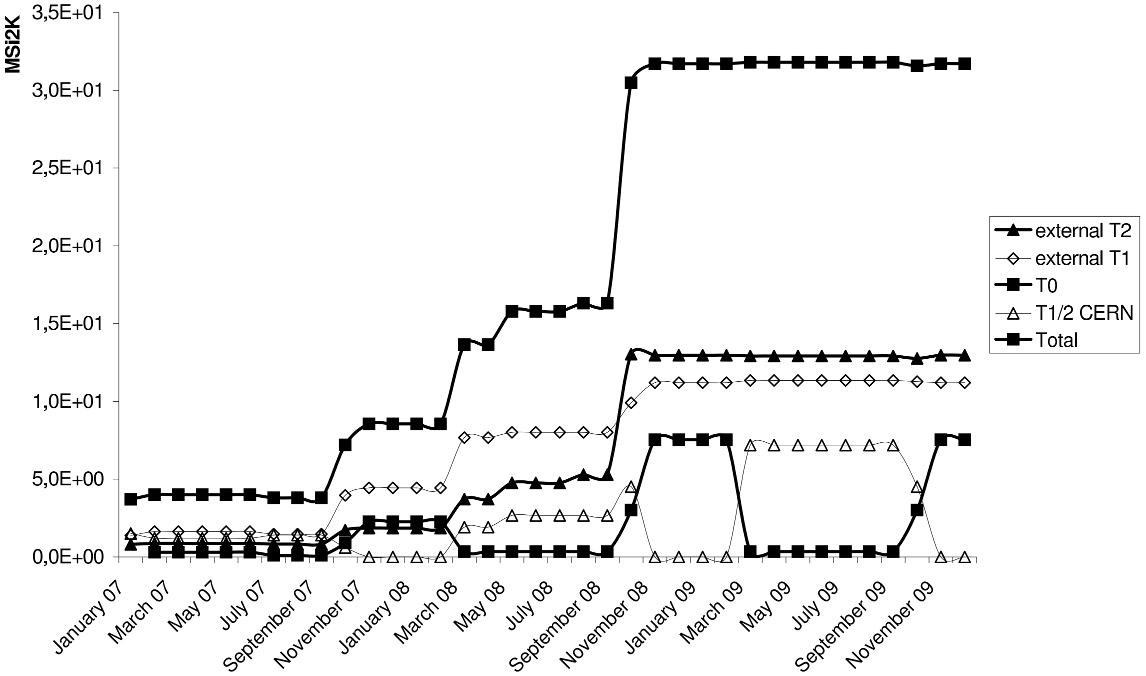

Figure 7.2: Profile of processing resources required at Tier 0, Tier 1s, and Tier 2s to process all ALICE data.

the processing resources required for the analysis we have considered that 200 physicists will exercise their algorithms many times on a small fraction of the data and produce physics results on a larger subset of the data. We furthermore assumed that the various analyses launched by the physics working groups are performed only once over the entire set of reconstructed data. Scheduled analysis, regrouping many analysis tasks, will require a larger amount of computing power per event than chaotic analysis. Individuals will also perform analyses using computing resources available at their home institutes. These resources are not accounted for in the present evaluation of the requirements.

The annual required resources to analyse real data are listed in Table 7.5.

From the scenario sketched in Fig. 7.1 on the preceding page, we deduce in Fig. 7.2 the profile of CPU resources required, including the ramp-up scenario.

The final values of resources required during a standard year of running at Tier 0, Tier 1s and Tier 2s are summarized in Table 7.8 on page 17.

7.5 Storage requirements

All data produced by the ALICE detector will be stored permanently for the duration of the experiment. This includes the raw data, a second copy of the raw data, one set of Monte Carlo data, reconstructed data from all reconstruction passes, analysis objects, calibration and condition data. A fraction of the produced data will be kept on short-term storage, providing rapid access to data frequently processed and for I/O dominated tasks. The media for permanent and transient storage will be decided according to the technology available. Currently the permanent storage medium is magnetic tape and the transient storage medium is disk. The ratio between disk and tape storage will change with time and will be dictated by the price–performance ratio. The parameters used to estimate the storage requirements are derived from the Data Challenge experience and are reported in Table 7.2 and Table 7.5.

7.5.1 Permanent storage

Raw data from pp and heavy-ion runs are copied onto permanent storage at Tier 0. They are further exported to the external Tier 1s, each one keeping a fraction of raw data on permanent storage, thus

Table 7.6: Size per event and per year of raw data produced by the ALICE experiment in pp and in Pb–Pb runs and by the processing of raw and Monte Carlo data.

	Real data (MB)				Monte Carlo data (MB)	
	Raw	ESD	AOD	Event catalogue	Raw	ESD
pp per event	1.0[a]	0.04	0.004	0.01	0.4	0.04
per year ($\times 10^9$)	1.1	0.18			0.4	0.1
Heavy-ion per event	12.5	2.5	0.25	0.01	300	2.5
per year ($\times 10^9$)	1.4	1.0			3.0	0.9

---

[a] It is assumed that five events are piled up.

providing a second copy. To cater for possible network interruptions, a disk buffer corresponding to the equivalent of 12 hours of heavy-ion data taking is foreseen in the DAQ cluster, and the equivalent of 2 days at the Tier 0. One copy of each reconstruction pass is stored on permanent storage at Tier 0 and external Tier 1s, most likely where they have been generated. One copy of Monte Carlo data is stored and distributed among the Tier 1s. The annual requirement for permanent storage in each Tier category is summarized in Table 7.8 on page 17. These values correspond to the needs during a standard year of running. Tier 2s are not required to provide permanent mass storage.

7.5.2 Transient storage

The requirements for transient storage on fast access media depend very much on the computing model, on the balance of available distributed processing resources, and on the network bandwidth and occupancy. In the absence of Grid simulation tools, our estimates on needed transient storage capacities are based on the requirement to store on disks data that will be frequently accessed primarily through non-scheduled tasks. They include:

- a fraction of the yearly raw data stored at Tier 0 (2%) and at each external Tier 1 (10%);
- one copy of the calibration and condition data at Tier 0 and each external Tier 1;
- two copies of the reconstructed data (ESD) of two reconstruction passes distributed onto Tier 0 and Tier 1s transient storage; a fraction of the ESD reside on disk at Tier 2.
- all analysis objects (AOD, event catalogue) distributed at Tier 2s and Tier 1s where they have been produced;
- one copy of generated Monte Carlo data are distributed in Tier 1s;
- two copies of reconstructed Monte Carlo distributed at Tier 2s where they have been produced.

The overall transient storage capacities required during a standard data taking year in each Tier category are listed in Table 7.8 on page 17.

7.6 Network

7.6.1 Tier 0

Data from the ALICE DAQ are transfered to Tier 0 at a rate of 1.25 GB/s during heavy-ion runs and 100 MB/s during pp runs. The DAQ farm provides for sufficient disk storage (50 TB) to buffer the equivalent of 12 hours of heavy-ion data taking. The maximum network bandwidth into Tier 0 is therefore

10 Gb/s during the heavy-ion run and 0.8 Gb/s during the pp run. The 10 Gb/s lines scheduled to be installed between the ALICE experiment and Tier 0 and exclusively dedicated to the raw data transfer will be sufficient.

Several kinds of transfer will contribute to the outgoing traffic from Tier 0 to external Tier 1s:

- the raw pp data export during the seven months of data taking leads to a rate of 55 MB/s;
- during the same time ESD data are exported at a rate of 3 MB/s;
- the raw heavy-ion data exportation during the month of data taking and the following four months of shutdown leads to a rate of 120 MB/s;
- during the four months of shutdown time ESD data are exported with a rate of 26 MB/s.

Data will be transferred simultaneously to every external Tier 1. This traffic results in a data transfer peak of 150 MB/s during the winter shutdown period. These values translate into a maximum outgoing network bandwidth of about 2 Gb/s (average over the year 1 Gb/s). A disk buffer of 200 TB is required at Tier 0 to store the equivalent of 48 hours of exported Pb–Pb RAW data.

7.6.2 Tier 1

The maximum incoming bandwidth required from Tier 0 into each Tier 1 is deduced from the discussion in the previous section to be 0.3 Gb/s assuming raw data are distributed in six Tier 1s. To the traffic from Tier 0 has to be added the traffic from Tier 2 exporting Monte Carlo data. Assuming that all Monte Carlo data are processed in Tier 2s and that there are on average 3.5 Tier 2s sending their locally produced data to a given Tier 1, this traffic amounts to 20 MB/s. The incoming bandwidth is therefore equal to about 2 Gb/s.

Tier 1 exports ESD data to Tier 2 for analysis. Assuming that all the analysis is done in Tier 2s and that each Tier 1 serves all the Tier 2s, the traffic from one Tier 1 to the Tier 2s amounts to 3 MB/s and requires a maximum bandwidth of about 0.03 Gb/s.

7.6.3 Tier 2

The incoming and outgoing bandwidth for Tier 2 is deduced from the discussion in the previous section to be 0.01 Gb/s and 0.6 Gb/s, respectively.

The requested network bandwidth for the various Tier categories is summarized in Table 7.8 on the next page. It should be noted that we assumed that the network is 100% efficient and operational 100% of the time. The network infrastructure should make provision for upgrades made possible by an evolving technology.

One element that makes this evaluation difficult is the duplication factor for disk files. It is clear that we cannot keep all the active data on disk without requiring an unreasonable amount of disk space. Data will be copied from Tier 1s permanent storage onto Tier 2s disk buffers. When the disk buffer is full, data will be removed and then copied again when needed. This may affect the network load between Tier 1s and Tier 2s and also between Tier 2s according to the end-user analysis pattern and the functionality of the middleware performing the file movement.

7.7 Ramp-up of resources

The total amount of resources required by ALICE for the production of Monte Carlo data and the processing of real data in a standard year of running, are summarized in Table 7.8 on the facing page. A small overall safety factor of 10% has been taken into account for all kinds of resources (CPU and storage). The staging of the offline resources (see Table 7.7) is dictated by the LHC heavy-ion schedule and is synchronized with the ramp-up of the DAQ online resources.

Table 7.7: Estimates for the computing resources ramp-up scenario. The timing of the resources increase is tuned to have the resources required in a given year available for the start of the heavy-ion run. The share of CERN CPU resources between Tier 0 and Tier 1/2 is describe in the text.

CPU (MSI2K)	2007	2008	2009	2010
CERN Tier 0	3.3	8.3	10.8	14.0
CERN Tier 1/2				
Ex Tier 1's	4.9	12.3	16.0	20.9
Ex Tier 2's	5.8	14.4	18.7	24.3
Total	14.0	35.0	45.5	59.2
Disk (TB)				
CERN Tier 0	95	238	309	402
CERN Tier 1/2	579	1447	1882	2446
Ex Tier 1's	2941	7353	9559	12426
Ex Tier 2's	2042	5106	6638	8629
Total	5658	14144	18387	23903
MSS (TB/year)				
CERN Tier 0	990	2475	3218	4183
CERN Tier 1	463	1158	1505	1957
Ex Tier 1's	2779	6947	9031	11740
Total	4232	10580	13754	17880

Before the first heavy-ion run, 20% of the nominal total resources are required to process pp data and to produce Monte Carlo data. ALICE foresees an early and short heavy-ion pilot run before the end of 2007 for which 40% of the resources required at CERN must be available already in mid 2007. 40% of the external resources must be available when the heavy-ion starts to take up the tasks of the CERN Tier 1 and Tier 2 which are not available during the heavy-ion run and during the first reconstruction pass (4 months after the heavy-ion run).

Table 7.8: Summary of the computing resources required by the ALICE computing model during a standard data-taking year (reference year 2008).

	CERN (Tier 0/1/2)	Tier 1s	Tier 2s	Total
CPU (MSI2k)	8.3	12.3	14.4	35.0
Transient storage (PB)	1.7	7.4	5.1	14.1
Permanent storage (PB/year)	3.6	7.0	–	10.6
Bandwidth in (Gb/s)	10	2	0.01	-
Bandwidth out (Gb/s)	1.2	0.03	0.6	-

The first full heavy-ion run is foreseen during the last quarter of 2008. Even though beams with reduced luminosity with respect to the nominal LHC luminosity might be available, data will be taken at the maximum rated allowed by the DAQ nominal bandwidth. In addition, even if the event size is smaller than the size used to estimate the resources required by the computing model, we will saturate the available DAQ bandwidth by tuning the triggers to increase the event rates, in such a way as to collect faster the statistics required for the ALICE physics goals. It would be unreasonable to limit the physics reach of ALICE during the startup years by delaying the installation at CERN of the full computing capacity. The rampup of the external resources must follow the rampup of resources at CERN to avoid to accumulate data to be processed. Delaying the installation of external resources would deprive ALICE from the needed capacity to react quickly to unexpected discoveries. Therefore, 100% of the required resources must be installed at CERN, as well as 100% of the external resources, already in 2008 for the first Pb–Pb run.

In 2009 and 2010, a 30% increase is requested in Tier 1 and Tier 2 resources to be able to cope with the reconstruction of the data of the running year and previous years. The computing resources required at CERN include Tier 0, Tier 1 and Tier 2 type resources. The first pass heavy-ion reconstruction will be performed during the four months after the heavy-ion run at Tier 0 and requests all resources (7.5 MSI2k) available at the CERN Tier 0. During the same time no Tier 1/2 resources are requested at CERN.

7.8 Summary

The computing resources requirements in each Tier category are summarized in Table 7.8. Our resources deployment scenario from 2007 to 2010 is summarized in Table 7.7.

The resources at CERN are shared between Tier 0, Tier 1, and Tier 2 services. The resources at CERN Tier 1 are sized larger than the resources requested in an average Tier 1, but are similar to those foreseen in the large Tier 1s, such as GridKa. The resources at CERN Tier 2 are of the size of the resources requested in an average Tier 2. The resources at Tier 0 have been evaluated to process the pp data quasi-on line, as specified by our Computing Model. Just after the end of the heavy-ion run, Tier 1 and Tier 2 services will be switched off at CERN and their resources will be exclusively devoted to Tier 0. In such a way, sufficient resources will be made available to perform the first-pass reconstruction of heavy-ion data within the four months following the heavy-ion run. This mode of operation (see Fig. 7.1) will leave a sufficiently large time buffer allowing us to analyse data and to define the running conditions for the next heavy-ion run.

Tier 1s perform scheduled tasks such as additional reconstruction passes and scheduled analysis. Tier 2s produce and process Monte Carlo data and perform unscheduled end-user analysis tasks. The processing outside CERN is balanced to cope with the absence of Tier 1 and Tier 2 resources at CERN during the first-pass reconstruction of heavy-ion data and to provide an uniform global CPU resource request in external Tier 1s and Tier 2s. The yearly requirement for permanent storage is not likely to change with time, whereas the disk storage capacity requirements might change depending on the performance of the upcoming mass storage systems on the one hand and on the performance of the Grid on the other.

References

Chapter 7

[1] LHC Computing Review, CERN/LHCC/2001-004.

[2] LHCC Review of Computing Resources for the LHC Experiments,
CERN-C-RRB-2005-006;
CERN-LHCC-2005-006;
LHCC-G-089.

[3] ALICE Collaboration, *Technical Design Report of Trigger, Data Acquisition, High-Level Trigger and Control System*, CERN/LHCC/2003–062.

[4] CERN/LHCC 2003-049, ALICE Physics Performance Report, Volume 1 (7 November 2003);
ALICE Collaboration: F. Carminati *et al.*, J. Phys. G: Nucl. Part. Phys. **30** (2004) 1517–1763.