Curriculum Vitæ of Patrick Valduriez

February 2025

Nationality: French

Born: 12 may 1955 in Calais, France

Prof. address: Inria at University of Montpellier
Bâtiment 5 - 860 rue de St Priest
34392 Montpellier Cedex 5 France

Tel: +33 6 70 01 31 97

e-mail: Patrick.Valduriez@inria.fr

http: http://www-sop.inria.fr/members/Patrick.Valduriez/

Position: Director of Research Emeritus Inria, Zenith team
Scientific advisor for International Relationships department at Inria
Scientific director of the Inria-Brasil strategic partnership
Chief scientific officer at LeanXcale
Collaborator at LNCC, Brazil

Education

1985: "Doctorat d'Etat" in computer science, University Pierre et Marie Curie (UPMC), Paris, France, 20 sept. 1985.
Title: ''Optimization of relational operators in database machines''.
Committee: M. Adiba (U. Grenoble), F. Bancilhon (MCC, Austin), D. DeWitt (U. Wisconsin), Ph. Flajolet (Inria), G. Gardarin (UPMC), C. Girault (UPMC).

1981: Ph.D. in computer science (Doctorat de 3ème Cycle), UPMC, 16 dec. 1981.
Title: ''Design and implementation of a join and sort processor for a database machine''.
Committee: M. Adiba (U. Grenoble), F. Bancilhon (U. Paris Sud), G. Gardarin (advisor, UPMC), C. Girault (UPMC), J. Lebihan (Inria).

1977: Master's degree in computer science (Maitrise d'Informatique), UPMC.

1975: Bachelor's degree in computer science (DUT Informatique), University of Lille.

Research Interests

Data science, data management, big data, databases, distributed and parallel data management, scientific data management, workflow management, data engineering, transaction management, replication, query processing, performance evaluation, cluster and cloud computing.

Ranking

Research.com (retrieved Feb, 2025): World Ranking=2680, National Ranking=47
Google Scholar (retrieved Feb. 2025): H-index=61, Citations=22,271
DBLP (retrieved Feb. 2025): Publications=328 (116 journals, 212 conferences)

Distinctions and Prizes

- AAIA Fellow, 2024.
- Innovation Prize, Inria & French Academy of Sciences, 2014.
- ACM Fellow, 2012.
- Who's Who France, 2010.
- Trustee Emeritus of the VLDB Endowment, 1998.
- IBM France Scientific Prize in Computer Science, 1993.

Best Paper Awards

- DEXA 2020 for "Distributed Caching of Scientific Workflows in Multisite Cloud".
- SBBD 2017 for "Spark Scalability Analysis in a Scientific Workflow".
- Colibri 2009 for "SARAVÁ: data sharing for online communities in P2P".
- VecPar 2008 for "High-performance Query Processing of a Real-world OLAP Database with ParGRES".
- SBAC 2005 for "Physical and Virtual Partitioning in OLAP Database Clusters" selected for International Journal of High-Performance Computing and Networking (special issue on best papers of SBAC 2005).
- VLDB 2000 for "PicoDBMS: Scaling down Database Techniques for the Smartcard".

Best Paper Nominations

- Selected among the (two) best papers of Globe2010 for Journal of Transactions on Large-Scale Data and Knowledge-Centered Systems (Springer) for "Continuous Timestamping for Efficient Replication Management in DHTs".
- Selected as second best paper of ISSRE2008 for Empiral Software Engineering Journal for "A Framework for Testing Peer-to-peer Systems".
- Nominated for best paper (top 4) of ESWC2008 for "Improving Interoperability Using Query Interpretation in Semantic Vector Spaces".
- Nominated for best paper (top 3) of SBBD2003 for "Cherry Picking: a semantic query processing strategy for the evaluation of expensive predicates."
- Selected for Network and Information Systems Journal (special issue of best papers of CIKM 1998) for "Dynamic Memory Allocation for Large Query Execution".

Professional Experience

- July 2023 - now: Collaborator at LNCC, Brazil, working on data science for environmental applications.
- April 2023 - now: Researcher in the Zenith team, working on data science for environmental applications.
- April 2019 - now: Chief Scientific Officer at LeanXcale, developing a new generation SQL system.
- Jan 2011 - March 2023: Manager of Zenith (Scientific Data Management) an Inria team (30 researchers and Ph.D students) at LIRMM, Montpellier. Zenith addressed the three main challenges of scientific data management: scale (big data, big applications); complexity (uncertain, multi-scale data with lots of dimensions); heterogeneity (in particular, data semantics heterogeneity). We propose innovative solutions with significant advantages in terms of scalability, functionality, ease of use, and performance. We designed and validated our solutions by working closely with scientific application partners (CIRAD, INRA, IRD, …) and industrial partners (Bull, Safran, LeanXcale, EDF, …).
- Sept 2002 - Dec 2010: Manager of Atlas (ATLantic dAta Systems), an Inria team (25 researchers and Ph.D students) on complex data management in distributed systems, Inria-Rennes Bretagne Atlantique and University of Nantes. Atlas's objective was to support complex data in the context of data grid computing, with large-scale distribution and parallelism. To scale up to very numbers

of nodes, we used a P2P approach. To ease large-scale data integration, we used a model-driven data engineering approach. Atlas has succeeded in a major technology transfer with two model-driven data engineering tools: the Atlas Transformation Language (ATL) and Atlas Model Weaver (AMW), now standard tools on the Eclipse platform as Open Source Software.

- Sept 2000 - August 2002: Professor of Computer Science at Université Pierre et Marie Curie – UPMC (Paris 6), on leave from Inria, and researcher at LIP6. Work on supporting autonomous databases and applications in large Linux-based clusters, addressing the issues of advanced transaction and replication models and load balancing to increase parallelism. In this context, I have started the RNTL funded project Leg@net (2002-2004) with Prologue Software and ASPLine to deal with legacy database application in an Application Service Provider context.

- April 1996 - March 2001: Director of Dyade (Bull-Inria R&D joint venture), 50 researchers and engineers, Rocquencourt. Dyade has associated Bull and Inria for 5 years to foster the development of basic technologies for advanced information systems. The joint venture had 50 people (Bull engineers and Inria researchers). Dyade has been considered a great success with several technology transfers in Bull products (Safetech, IPv6, OpenMaster, Netwall) and four startups: TrustedLogic (formal methods for security) in 1998, Kelkoo (now number one Internet buying guide in Europe) in 1999, Jalios (collaborative editing) in 2000 and Scalagent (software agents) in 2001. Kelkoo uses the Disco technology, which I developed in the Rodin project.

- June 1989 - June 1999: Manager of the Rodin project team, 20 researchers and Ph.D. ,, Inria Rocquencourt. The general objective of the Rodin group was to design and experiment with advanced database techniques. The Rodin group has been internationally recognized for major contributions in query processing for distributed and parallel systems, and distributed object management. The Rodin group achieved several technology transfers to industry, including the Distributed Information Search Component (Disco) which Inria licensed to Kelkoo in 1999.

- January 1987 - May 1989: Team leader, 6 people, Microelectronics and Computer Technology Corp. (MCC), Austin, Texas. Responsible for the design and implementation of FAD, a database programming language for a parallel computer (Bubba), and its compiler/optimizer. A working prototype was available in early 1989 for a 40 node Flexible parallel computer, and eventually transferred to MCC shareholders. In particular, it was the basis for NCR's Navigator product (developed with Sybase) in 1994.

- December 1984 - January 1987: Researcher, MCC, Austin, Texas. Design and implementation of a parallel database system prototype on a network of Sun workstations. The prototype includes the concept of Join Index which I invented to speed up complex query processing in 1986. Since then, all the major DBMS vendors have implemented join indices.

- January 1982 - December 1984: Researcher, Inria, Rocquencourt. Leader of a Query Processor team with 4 Ph.D. students. Design and implementation of a query processor for a relational DBMS on Multics (DPS8) and Unix (SM90). The technology was transferred into the Sabrina DBMS product to Infosys, a startup created by Inria researchers.

- October 1979 - December 1981: Doctoral Researcher, Inria, Rocquencourt. Design and implementation of parallel algorithms for the execution of relational operators in databases. One of the first demonstrations of the superiority of hashing for parallel query processing.

- December 1977 - September 1979: Programmer (Volontaire du Service National Actif (VSNA), Ministère de la Coopération), Office Béninois d'Informatique, Cotonou, Bénin. Design and

implementation of large banking and statistical applications on a mainframe system (Honeywell-Bull HB68 runing the GCOS operating system).

Software and Technology Transfer

Major Transfers in Commercial Products

- Triton End-to-end Graph Mapper, Beepeers, 2017. A server for managing graph data and applications for mobile social networks. Developped in the context of the Triton I-lab and transferred to the Beepeers startup (www.beepeers.com), with a participation of Inria in the company. I participated to the design (contribution 20%).

- CloudMdsQL compiler, LeanXcale, 2017. CloudMdsQL is a new query language for multistore systems. We developed CloudMdsQL and its compiler/optimizer in the context of the CoherentPaaS European project and transferred it to the LeanXcale company (www.leanxcale.com) for its data analytics engine in 2016. I led the work (contribution 40%) with one engineer, one postdoc and one PhD. Student.

- WebSmatch, DataPublica, 2014. WebSmatch is a flexible, open environment for discovering and matching complex schemas from many heterogeneous data sources over the Web. It has been developed in the context of the Xdata PIA and used by the Datapublica startup (acquired by SideTrade (www.sidetrade.com) in 2018) as a basis for its Radar product. I managed the WebSmatch project with 2 researchers from University of Montpellier and one engineer. My contribution is 40%.

- DISCO (Distributed Information Search Component), Kelkoo 1999. Disco is a data integration system developed in the Rodin group (in the context of Dyade with Bull) between 1995 and 1999. I led the Disco team (A. Tomasic – architect, H. Naacke and three Ph.D. students), participated in the design of Disco and transferred it to Kelkoo (now number one Internet buying guide in Europe and a Yahoo company) in 1999. We registered Disco to APP (Agence pour la Protection des Programmes) in 1999; my contribution to the product is 20%. According to Inria Transfer's head Laurent Kott, this technology transfer was exemplary.

- Garbage collector, O2 Technology, 1997. With M. Skubiszewski (then in my Rodin group), we designed a concurrent garbage collector for the O2 DBMS in 1997. The technology has been transferred into the product by M. Skubiszewski. My contribution is about 20%.

- Navigator, NCR 1989. Navigator is a parallel query processor commercialized by NCR with Sybase in 1994. It was later on integrated in NCR's flagship product Teradata (the best selling DBMS product for very large databases). NCR has used the technology which I developed in the FAD query processor at MCC (my contribution = 40%). The technology transfer was done by Brian Hart, a member of my FAD team at MCC.

- Sabrina, Infosys 1984. Sabrina is a relational DBMS which we developed in the SABRE group (headed by G. Gardarin) at Inria-Rocquencourt. It has been transferred to Infosys, an Inria startup, and sold at a hundred copies. I developed the query processor and part of the execution system (relational operator algorithms), i.e., about 20% of the product.

Major Open Source Software products

- DfAnalyzer : https://github.com/vssousa/dfanalyzer-spark. Since 2018. A tool for monitoring, debugging, steering, and analysis of dataflows while being generated by scientific applications. It works by capturing strategic domain data, registering provenance and execution data to enable

queries at runtime. DfAnalyzer provides lightweight dataflow monitoring components to be invoked by high performance applications. It can be plugged in scripts, or Spark applications, in the same way users already plug visualization library components.

- SAVIME (Simulation And Visualization IN-Memory): https://github.com/hllustosa/Savime. Since 2017. A multi-dimensional array DBMS for scientific applications. SAVIME supports a novel data model called TARS (Typed ARray Schema), which extends the basic array data model with typed arrays. In TARS, the support of application dependent data characteristics, such as data visualization and UQ computation, is provided through the definition of TAR objects, ready to be manipulated by TAR operators. This approach provides much flexibility for capturing internal data layouts through mapping functions, which makes data ingestion independent of how simulation data has been produced, thus minimizing ingestion time. Deveopped in the context of the SciDISC associated team with LNCC, Rio de Janeiro.

- Hadoop_g5k: https://github.com/mliroz/hadoop_g5k/wiki. Since 2014. Apache Hadoop provides an open-source framework for reliable, scalable, parallel computing. It can be deployed and used in large-scale platforms such as Grid 5000. However, its configuration and managementis very difficult, especially under the dynamic nature of clusters. Therefore, we built Hadoop_g5k (Hadoop easy deployment in clusters), a tool that makes it easier to manage Hadoop clusters and prepare reproducible experiments. Hadoop_g5k offers a set of scripts to be used in command-line interfaces and a Python interface. It is actually used by Grid5000 users, and helps them saving much time when doing their experiments with MapReduce.

- CloudMdsQL Compiler: http://cloudmdsql.gforge.inria.fr. Since 2013. The CloudMdsQL (Cloud Multi-datastore Query Language) compiler transforms queries expressed in a common SQL-like query language into an optimized query execution plan to be executed over multiple cloud data stores (SQL, NoSQL, HDFS, etc.) through a query engine. The compiler/optimizer is implemented in C++ and uses the Boost.Spirit framework for parsing context-free grammars. CloudMdsQL has been validated on relational, document and graph data stores in the context of the CoherentPaaS European project, and transfeered to the LeanXcale startup.

- SciFloware: http://www-sop.inria.fr/members/Didier.Parigot/pmwiki/Scifloware. Since 2013. SciFloware is a middleware for the execution of scientific workflows in a distributed and parallel way. It capitalizes on our experience with SON and an innovative algebraic approach to the management of scientific workflows. SciFloware provides a development environment and a runtime environment for scientific workflows, interoperable with existingsystems. We validate SciFloware with workflows for analyzing biological data provided by our partners CIRAD, INRA and IRD.

- WebSmatch (Web Schema Matching): http://websmatch.gforge.inria.fr/. 2011-2014. WebSmatch is a flexible, open environment for discovering and matching complex schemas from many heterogeneous data sources over the Web. It provides three basic functions: (1) metadata extraction from data sources; (2) schema matching (both 2-way and n-way schema matching), (3) schema clustering to group similar schemas together. WebSmatch is being delivered through Web services, to be used directly by data integrators or other tools, with RIA clients. Implemented in Java, delivered as Open Source Software (under LGPL) and protected by a deposit at APP (Agence de Protection des Programmes). WebSmatch has been used by Datapublica and CIRAD to integrate public data sources.

- P2Prec (P2P recommendation service): http://p2prec.gforge.inria.fr. 2010-2013. P2Prec is a recommendation service for P2P content sharing systems that exploits users social data. To

manage users social data, we rely on Friend-Of-A-Friend (FOAF) descriptions. P2Prec has a hybrid P2P architecture to work on top of any P2P content sharing system. It combines efficient DHT indexing to manage the users FOAF files with gossip robustness to disseminate the topics of expertise between friends.

- ProbDB (Probabilistic Database): http://probdb.gforge.inria.fr. 2010-2013. ProbDB is a probabilistic data management system to manage uncertain data on top of relational DBMSs. One of the main features of the prototype is its portability; that means with a minimum effort it can be implemented over any DBMS. In ProbDB, we take advantage of the functionalities provided by almost all DBMSs, particularly the query processing functions. It is implemented in Java on top of PostgreSQL.

- ATL (Atlas Transformation Language): 2004-2006. ATL is a transformation-based model management framework, with metadata management and data mapping as the main applications. It comes with a library of more than 100 transformation components. ATL has been registered in 2004 (by Inria, TNI-Software and University of Nantes) to the APP. The main authors of ATL are J. Bézivin (20%), F. Jouault (60%) and myself (10%). ATL is released as Open Source Software under the Eclipse Public Licence and available as an Eclipse plugin. The average number of downloads is 675 per month. There is now an active community of more than 100 user sites, including research labs and major companies (Airbus, NASA, Ilog, Sodius, TNI, etc.). ATL has been selected by two competitivity centers in France: System@tic and TopCased. In early 2007, ATL was recognized a standard Eclipse component for model transformation.

- AMW (Atlas Model Weaver): 2004-2006. AMW is a component-based platform for model weaving, i.e. establishing and managing correspondences between models, with data integration as main application. AMW has been registered in 2005 (by Inria and University of Nantes) to the APP. The authors of ATL are J. Bézivin (20%), F. Jouault (10%), M. Didonet del Fabro (20%) and myself (20%). AMW is released as Open Source Software under the Eclipse Public Licence and available as an Eclipse plugin. The average number of downloads is 210 per month. AMW is being used by more than 40 user sites, including research labs and major companies (NASA, BAE, Versata, Obeo, etc.).

- APPA (Atlas P2P Architecture): 2006-2009. P2P data management system that provides scalability, availability and performance for applications that deal with semantically rich data (XML, relational, etc.). APPA provides advanced services such as queries, replication, privacy, and testing. It has been implemented on top on various P2P networks such as JXTA, OpenChord and Pastry and tested on GRID5000 and PlanetLab. The current services of APPA are (see below): KTS, SbQA, P2P-LTR and PeerUnit. A privacy service (Priserv) is being implemented. The APPA services have been used in several projects: Strep Grid4All, ANR RNTL Xwiki Concerto and ANR VERSO DataRing. My contribution to APPA is about 20%.

- KTS (Key-based Timestamp Service): 2007-2009. A distributed service to manage timestamps in DHTs. It is useful to solve various DHT problems which need a total order on operations performed on each data. KTS has been initially proposed to support data currency in DHTs, i.e. the ability to return a current replica in a DHT despite peers leaving the network or concurrent updates. Experimental validation has shown that KTS incurs very little overhead in terms of communication cost. KTS is the basis for the P2P-LTR service. It has been implemented in Java on top of OpenChord.

- P2P-LTR (P2P Logging and Timestamping for Reconciliation): 2007-2009. Provides two major functions: logging of user actions in a DHT and continuous, distributed timestamping of these

actions. This is useful to perform reconciliation of replicated data. P2P-LTR extends KTS with continuous timestamping and logging of actions. To perform reconciliation using P2P-LTR, we use a simple reconciliation algorithm based on operational transforms, called SB, from the ECOO team at LORIA and readily available as Open Source Software. P2P-LTR has been implemented in Java on top of OpenChord. It has been validated in the Strep Grid4All and RNTL Xwiki Concerto projects to perform reconciliation of replicated documents in a P2P wiki system.

- SbQA (Satisfaction-based Query Allocation Framework): 2007-2009. Query allocation framework for distributed environments where consumers and providers are autonomous and have special interests towards providers and queries, respectively. We experimentally demonstrated that it ensures good system performances while satisfying consumers and providers. Hence, SbQA can scale-up in these environments by preserving the total system capacity, i.e. the aggregate capacity of all providers. SbQA is used in the Strep Grid4All project as the basis to perform selection of services proposed by market-places as well as altruist contributors. SbQA is implemented in Java.

- PeerUnit (Peer-to-Peer Tester): 2006-2009. Testing framework for P2P systems. It is useful to developers who want to implement unit tests for a Java P2P system. The framework is based on two original aspects: (i) the individual control of peers volatility and (ii) a distributed testing architecture to cope with large numbers of peers. A distributed component, the tester, executes on peers, and controls their execution and their volatility, making them leave and join the system at any time, according to the needs of a test. Furthermore, testers communicate with each other across a balanced tree (B-Tree) structure to avoid using a centralized testing coordination. Peerunit is implemented in Java and has been validated on two popular open-source P2P systems (FreePastry and OpenChord).

- RepDB*: 2002-2005. RepDB* is a data management component for replicating autonomous databases or data sources in a cluster system. It is implemented in Java on Linux and supports various DBMS: Oracle, PostGreSQL and BerkeleyDB. It has been validated on GRID5000 with a 64-node cluster. At the end of 2004, RepDB* was registered (by Inria and the University of Nantes) to the APP. The authors of RepDB* are C. Coulon (25%), G. Gaumer (10%), E. Pacitti (50%) and myself (15%). RepDB* was released as Open Source Software under the GPL licence. There has been more than 2000 downloads in the first months of release.

Main Other Research prototypes

- Triton End-to-end Graph Mapper: https://hal-lirmm.ccsd.cnrs.fr/lirmm-01620239. Since 2017. A server for managing graph data and applications for mobile social networks. The server is built on top of the OrientDB graph database system and a distributed middleware. It provides an End-to-end Graph Mapper (EGM) for modeling the whole application as (i) a set of graphs representing the business data, the in-memory data structure maintained by the application and the user interface (tree of graphical components), and (ii) a set of standardized mapping operators that maps these graphs with each other. In the context of the Triton I-lab, the technology has been transferred to the Beepeers startup.

- PicoDBMS: 1999-2001. PicoDBMS is a lightweight DBMS embedded in a smartcard for secured folders. PicoDBMS is now developped in the Smis group at Inria-Rocquencourt. I participated in the design of the first version of PicoDBMS with L. Bouganim and P. Pucheral in 1999-2000 (my contribution = 25%). We obtained the VLDB2000 best paper award for the design of PicoDBMS. We also demonstrated the first prototype in Javacard at VLDB2001.

Professional Activities

Management of research and R&D

- Director of the Inria-Brasil strategic partnership, since jan 2020.
- Manager of the Zenith Inria team (30 researchers) at LIRMM, jan 2011-march 2023.
- Head of the workpackage on biology data integration at Institut de Biologie Computationnelle (IBC), Montpellier, 2012-2018.
- Head of the scientific data group at Labex NUMEV, Montpellier, 2010-2014.
- Head of the ModSiCS2020 (Data, Models and Theories for Complex Systems in 2020) Working Group, University Montpellier 2, 2010-2011.
- Manager of the Atlas Inria team (25 researchers), sept. 2002-dec. 2010.
- Director of European programs and international affairs, LINA, 2003-2008.
- Director of GIE Dyade between Bull and Inria (50 researchers and engineers), 1996-2001.
- Manager of the Rodin Inria project team (25 researchers) at Inria-Rocquencourt, 1989-1999.
- Leader of the FAD team (6 researchers) at MCC, Austin, Texas, 1987-1989.
- Leader of the Query Processor team (4 Ph.D. students) in the SABRE group, Inria-Rocquencourt 1982-1984.

Administration of research

- President of the BDA steering committee (2017-2020).
- Agence Nationale de la Recherche (ANR) Member of the Scientific Committee: Défi 7 Information and communication society, 2014-2015.
- Inria International Affairs - Direction des Relations Internationales (DRI): scientific manager for the Latin America zone, since January 2014.
- Member of ACM SIGMOD steering committee, 2001-2004.
- Elected trustee of the VLDB Endowment, 1993-1997.
- Member of the ACI GRID scientific committee, Ministry of Research (2002-2004).
- Director of the GDR BD, Ministry of Research (1989-1993).
- President of the Database technical committee of AFCET (1989-1993).

Associate editor of international journals

- The VLDB Journal, Springer (1998-2001, 2012-2015).
- Transactions on Database Systems, Association for Computing Machinery (1993-1998).
- Int. Journal on Intelligent and Cooperative Database Systems, World Scientific (1998-2002).
- Proceedings of the VLDB Endowment (2009-2012).
- Distributed and Parallel Database Systems, Kluwer Academic Publishers (since 1998).
- Journal of Information and Data Management (2010-2016).
- Internet and Databases: Web Information Systems, Kluwer Academic Publishers (2002-2010).
- Information System Engineering, Hermes, 2000-2004.
- Network and Information Systems Journal, Hermes (Editor-in-Chief, 1997-2000).

General chair of international conferences

- LaDAS@VLDB workshop, Rio de Janeiro, 2018.
- Int. Conf. on Very Large Data Bases (VLDB), Lyon, 2009 (750 participants).
- Extending Database Technology (EDBT), Nantes, 2008 (350 participants).
- High-performance Data Management in GRIDs (HPDgrid), Rio de Janeiro, 2006 (50 participants).
- ACM-SIGMOD Int. Conf., Paris, 2004 (650 participants).
- International Workshop on Information Integration on the Web, Rio de Janeiro, 2001 (100 participants).

Program committees of international conferences

- ACM-SIGMOD Int. Conf.: 1986, 1988, 1990, 1991, 1995, 1997 (PC chair), 1998, 2001, 2004, 2007, 2008, 2011, 2014 (area chair).
- Int. Conf. on VLDB: 1982, 1988, 1989, 1993, 1994, 1995, 1998 (industrial program chair), 1999 (European PC chair), 2000, 2002, 2003 (industrial program), 2005, 2006, 2007, 2008, 2011, 2015 (sponsor co-chair), 2018 (sponsor co-chair).
- IEEE Int. Conf. on Data Engineering: 1986, 1987, 1988, 1989, 1992, 2003, 2007 (ind. Prog. Chair), 2011, 2014, 2015 (DESWeb), 2016 (area chair), 2018 (sponsor co-chair).
- Int. Conf. on Extending Data Base Technology: 1990, 1992, 1996, 2010.
- IEEE Int. Conf. on Distributed Computing Systems: 2007 (data management track chair).
- Euro-par: vice-president (2005), president (2006) of Distributed and Parallel Databases track.
- ADBIS East-European Conference on Advances in Databases and Information Systems 2015: P. Valduriez (PC chair).
- Int. Conf. on Web Information System Engineering (WISE), 2003.
- Int. Conf. on Distributed and Parallel Information Systems: 1991, 1993 (PC chair).
- Journées Bases de Données Avancées (BDA) : 1990 (PC chair), 1992 (conference chair), 1994, 2001, 2003, 2007, …

Expert for international and national organizations

- High Council for Evaluation of Research and Higher Education (HCERES)
- Agence Nationale de la Recherche (ANR)
- Natural Sciences and Engineering Research Council of Canada (NSERC)
- National Science Foundation - USA (NSF)
- Commission of the European Communities (CEC)
- United Nations University (UNU)
- Canadian Research Council
- North Atlantic Treaty Organization (NATO)
- German Israeli Foundation for Scientific Research and Development.
- Centre National de la Recherche Scientifique (CNRS).
- Agence Nationale pour la Valorisation de la Recherche (ANVAR)
- Réseau National de la Recherche en Télécommunications (RNRT)

Consultant in IT for major companies

- USA: MCC, HP Labs, Lucent Bell Labs, Microsoft, NERA, LECG, Birdstep
- Europe: ESA, Eurocontrol, NATO, Shell, Ask
- France: Bull, Capgemini, Matra Marconi Space, Schlumberger, Sodifrance, Orsys, Murex, IRSN

Teaching

- The textbook Principles of Distributed Database Systems, co-authored with professor Tamer Özsu, U. Waterloo, Canada, published by Prentice Hall in 1991 and 1999 (2nd edition), by Springer in 2011 (3rd edition – 850 pages), by Tsinghua University Press in 2013 (Chinese Edition), and by Springer (4th edition – 674 pages) in 2020 has become the standard book for teaching distributed databases in universities all over the world at the graduate and undergraduate levels. Our Web site features course material, exercises, and direct communication with instructors.

- Teaching while professor at UPMC, Paris, 2000-2002 (2 years, 192 hours/year)
  Key-words: Databases, Object Technology, Distributed Databases

Public: Bs. (Licence 3) and Ms. (Master 1, Master 2)

- Teaching for professional schools in 1980-2023 (about 50 hours/year)
  Key-words: Databases, Distributed Information Systems, XML
  Public: Engineers, project managers
  Schools: Capgemini Institut, Orsys Formation

- Teaching since 2009 (about 20 hours/year)
  Key-words: Distributed and parallel data management
  Public: Ms. (Master 2)
  Schools: U. Montpellier

- Teaching in 1980 - 2009 (about 30 hours/year)
  Key-words: Databases, Object Technology, Distributed Systems, Parallel Systems
  Public: Ms. (Master 1, Master 2)
  Schools: UPMC, UVSQ, ESSEC, ENST, EMN, U. Nantes

Invited talks

1. "LifeSWS: motivations and architecture" on 4 September 2023, U. Waterloo, Canada.
2. "Data Science and Innovation" on 12 April 2023 at Inria-Brasil Workshop, Petropolis, Rio de Janeiro, Brazil.
3. "Innovation: startup strategies" on 19 May 2021 at Stratégies d'Innovation – 12ème edition, Marcusevans France (virtual).
4. "Distributed Database Systems: the case for NewSQL" on 19 November 2020 at CWI Lectures on Database Research, Amsterdam, Netherlands (virtual).
5. "Data-intensive Science" on 6 November 2019 at HPC - AI - BigData Convergence Days (Conv'2019), EDF Lab Paris-Saclay.
6. "The Case for Hybrid Transaction Analytical Processing" on 25 April 2019 at IBM Research Brazil, Rio de Janeiro, Brazil.
7. "Comment réussir sa thèse : bonnes pratiques et retours d'expériences" on 11 december 2018 at Journées des Doctorants du LIRIS, Lyon.
8. "Blockchain 2.0: opportunities and risks". 25 october 2018 at BDA Conf., Bucharest, Romania, 7 november 2018; AMSUD workshop, Bucaramanga, Colombia; 19 December 2018, Colloquium of PESC/COPPE, UFRJ, Rio de Janeiro, Brazil.
9. "Blockchain 2.0: opportunities and risks in developing countries". 2nd EAI International Conference on Emerging Technologies for Developing Countries (Africatek), Cotonou, Bénin, 29 May 2018.
10. "Data Science and Science Data", 1 December 2017, #Digitag Institut de Convergence on numerical agriculture, Montpellier.
11. "An Overview of Polystores" on 11 October 2017 at XLDB 2017, Clermont-Ferrand, France.
12. "Data Science: opportunities and risks" on 26 July 2017 at Colloquium of PESC/COPPE, UFRJ, Rio de Janeiro, Brazil.
13. "The CloudMdsQL Multistore System" on 2 May 2017 at Data Systems Group Seminar Series, University of Waterloo, Canada, and on 4 May 2017 at McGill University, Montréal, Canada.
14. "From Databases to Data Science: impact on information systems", 15 September 2016 at Junior Conference on Data Science and Engineering, Paris Saclay, France.
15. "Design and Implementation of the CloudMdsQL Multistore System" 24 April 2016 at 6th International Conference on Cloud Computing and Services Science (CLOSER 2016), Roma, Italy.
16. "Data Science: opportunities and risks", 11 April 2016 at Journée "La Science des données à l'IRIT", Toulouse, France.
17. "Big Data Management in Zenith", Fundação Getulio Vargas, Rio de Janeiro, 22 December 2015.

18. "Cloud & Big Data: opportunities and risks for developing countries", 15 December 2015 at 7th EAI International Conference on e-Infrastructure and e-Services for Developing Countries (AFRICOMM), Cotonou, Bénin.
19. "Data-intensive HPC: opportunities and challenges", Workshop on Big Data and Extreme-scale Computing (BDEC), Barcelona, 28-20 January 2015.
20. "CloudMdsQL: Querying Heterogeneous Cloud Data Stores with a Common Language", Inria-Technicolor Workshop on Storage and Processing of Big Data, Rennes, 3-4 December 2014.
21. "CloudMdsQL: Querying Heterogeneous Cloud Data Stores with a Common Language", Workshop on DB Consistency in the Cloud, LIP6, Paris, 15-16 September 2014.
22. "Cloud & Bigdata: opportunities and risks", New Approaches to Economic Challenges - OECD-ECLAC Workshop, Paris, 19 May 2014.
23. "Indexing and Processing Big Data", Colloquium Indexing for Scientific Big data, CNRS, Paris, 15 January 2014.
24. "Cloud & Bigdata: a perfect marriage?" Colloquium Jacques Morgenstern, Sophia Antipolis, 7 November 2013.
25. "Parallel Techniques for Big Data", Workshop Mastodons@Montpellier: Gestion de Données à Grande Echelle en Science de la Vie, Montpellier, 7 December 2012.
26. "Principles of distributed Data Management in 2020?" Int. Conf. on Databases and Expert Systems Applications (DEXA), Toulouse, France, 2011.
27. "Cloud Data Management", DNAC Congress, Paris, 2010.
28. "L'internet social (ou Web 2.0): opportunités, impact et défis", Séminaire Inter-universitaire de Philosophie et d'Histoire des Sciences (HiPhis), Montpellier, 2010. http://www.paroledechercheurs.net/spip.php?article768.
29. "Data Management in Large-Scale P2P Systems", Int. Conf. on High-performance Computing for Computational Science (Vec-Par), Valence, 2004.
30. "From XML to Universal Web Data Exchange", UFRJ, Rio de Janeiro, 2000.
31. "Enterprise Application Servers", Embratel Conf., Sao Paulo, 1999.
32. "Intranet et Systèmes d'Information Distribués", Conf. Rubis, Lausanne, 1998
33. "The Dyade R&D joint venture", Bell Labs, USA, 1998.
34. "Intranet et Groupware", Conf. de l'association des informaticiens du Magreb, Rabbat, 1997.
35. "Parallel and Distributed Database Systems", Int. Conf. on Parallel and Distributed Computing Systems, Dijon, 1996.
36. "Databases and Web", Conf. Africaine sur la Recherche en Informatique (CARI), Libreville, 1996.
37. "Query Processing in the IDEA System", Colloquium Series, U. d'Alberta, Edmonton, 1995.
38. "Databases and parallelism", Conf. ORAP, Paris, 1995.
39. "Rule-based Query Processing in the IDEA System", Int. Symp. on Advanced Database Technologies and their Integration, Nara, Japon, 1994.
40. "Advances in Distributed Database Systems", Pre-VLDB Int. Workshop, Rio de Janeiro, 1994.
41. "Parallel Database Systems : the case for shared-something", IEEE Int. Conf. on Data Engineering (ICDE), Vienne, 1993.
42. "Object Database Systems", Asian Institute of Technology, Bangkok, 1993.
43. "Query Processing in the EDS Parallel Database System", Brazilian Symp. on Databases, Rio de Janeiro, 1990.
44. "Parallel Recursive Query Processing in a Shared Nothing Data Server", Journées BDA, Benodet, 1988.

Others: U. of Washington (2007), EPFL (2005), U. Waterloo (2005), Bell Labs. (1998, 1997), HP Labs. (1995), IBM Almaden Research (1993), U. Maryland (1993), GTE Labs. (1993), UT Austin (1993), MCC (1993), etc.

Tutorials

1. "Principles of Distributed Database Systems: spotlight on NewSQL" on 29 september 2020 at Brazilian Symposium on Databases (SBBD), Rio de Janeiro, Brazil (virtual).
2. "NewSQL: principles, systems and current trends" on 19 december 2019 at IEEE Bigdata Conference, Los Angeles, USA.
3. "Distributed and Parallel Data Processing", CEA-EDF-Inria Summer School on Big Data Analytics, 16-20 june, Cadarache, France, 2014.
4. "Parallel Techniques for Big Data", Journées Bases de Données Avancées (BDA 2013), Nantes, 22 october 2013.
5. "Parallel Techniques for Big Data" Computational Biology Institute, Montpellier, 22 March 2013.
6. "XML and Data Management", Consist Days, Sao Paulo, Brazil, 2002.
7. "Information Integration with XML" Journées BDA, Agadir, Marocco, 2001.
8. "XML and software engineering", ICSSEA, Paris, 2000.
9. "Intranet and Distributed Information Systems", Journées BDA, Tunis, 1998.
10. "Object Databases", DBWorld, Londres, 1995.
11. "Distributed Databases", Europ.Sem. on Advances in Distributed Systems, Alpe d'Huez, 1995.
12. "Parallel Databases", Journées BDA, Clermont-Ferrand, 1994.
13. "Object Database Systems", TOOLS Europe, Versailles, 1993.
14. "Distributed Databases", Journées BD3, Trégastel, 1992.
15. "Object Databases", Int. Conf. on VLDB, Vancouver, 1992.
16. "Object Technology", Int. Conf. on DKSME, Lyon, 1992.
17. "Object Database Systems", Schlumberger Int. Conf., Versailles, 1991.
18. "Object and Distributed Database Systems", Database school, La Trobe Univ., Melbourne, 1991.
19. "Interoperable Database Systems", EDBT School, Algheiro, 1991.
20. "Parallel and Distributed Database Systems", Int. Conf. on VLDB, Brisbane, 1990.
21. "Parallel Databases", Journées BDA, Montpellier, 1990.

Research Grants

European grants

- H2020 CloudDbAppliance European Project with Bull/Atos (leader), Inria Zenith, U. Madrid, INESC and the companies LeanXcale, QuartetFS, Nordea, BTO, H3G, IKEA, CloudBiz, and Singular Logic 2016-2019 (3 years, 5M€), 500K€.
- H2020 HPC4E Brazil-European Project by Barcelona Supercomputing Center (BSC) with Inria, Lancaster University, Repsol, and Total in Europe, and COPPE/Universidade Federal de Rio de Janeiro, LNCC, Instituto Tecnológico de Aeronáutica (ITA), Universidade Federal do Rio Grande do Sul, Universidade Federal de Pernambuco and PETROBRAS in Brazil, 2015-2017 (2 years), 120K€.
- IP CoherentPaaS European Project with U. Madrid (leader), Inria Zenith, FORTH, ICCS, INESC and the companies MonetDB, QuartetFS, Sparsity, Neurocom, Portugal Telecom, 2013-2016 (3 years, 5M€), 500K€.
- STREP Grid4All European Project with France Telecom R&D, Kungliga Tekniska Hoegskolan (Sueden), SICS (Sueden), ICCS (Greece), University of Piraeus Research Center (Greece), Universitat Politècnica de Catalunya (Spain), Rededia S.L. (Spain), 2006-2008 (3 years) 200 K€.
- IP ModelPlex European Project with Thales Research TRT, Thales Information Systems, Softeam, IBM Belgium, Fraunhofer Gesellschaft, WesterGeco, IBM Israël, University of York, Universidad Politecnica de Madrid, UPMC, Fundacion European Software Institute, Adaptive Limited, XJTech, Xactium, Imbus AG, SINTEF, Sodifrance-MIA Software, Technical university Braunschweig, University Dresden, Telefonica, CEA, SAP., 2006-2009 (3 years), 335 K€.
- IP Modelware European Project with Thales, Softeam, IBM UK, As Aprote, Enabler Informatica, Telenor Communication, France Telecom, Logon Technology Transfer, Fraunhofer Gesellschaft, WesterGeco, IBM Israël, University of York, Universidad Politecnica de Madrid, UPMC, Fundacion European Software Institute, Adaptive Limited, Imbus AG, SINTEF, Interactive Objects Software, 2004-2006, 330 K€.

- ESPRIT Long Term Research DWQ (Data Warehouse Quality) with universities of Athens, Rome, and Aarhen, 1996-1999, 200 K€.
- ESPRIT 3 A Middleware to Interoperate with Multiple data Sources (Miroweb) with Bull, Ibermatica, OSIS, GMD, 1997-1999 (2 years), 200 K€.
- ESPRIT 3 Intelligent Database Environment for Advanced Applications (IDEA) with Bull, ICL, ECRC, 1992-1994 (3 years), 500 K€.
- European Declarative System (EDS) with Bull, ICL, Siemens, ECRC, Infosys, ESPRIT 2 EP2025, 1989-1992 (4 years), 600 K€.
- Fully Integrated Data Environment (FIDE) with U. Glasgow, St Andrews, Pisa, Hamburg, ESPRIT 3 BRA 6309, 1992-1995 (3 years), 150 K€.

Industrial grants
- Data analysis of big time series, SAFRAN, since 2017, 100K€.
- Advanced data storage and processing for cloud workflows, Microsoft Research, 2013-2017, 120K€.
- Privacy preserving data mining in P2P networks, EDF, 2013-2014, 120 K€.
- Model engineering and data management, Microsoft Research, 2003-2006, 80 K€.
- Model management; IBM/Eclipse, 2004-2005, 20 K€.
- Study of access to distributed catalogs, Alcatel Space Industries, 2000, 15 K€..
- Study of Nutrition database techniques, Nutrimate, 1999 (3 months), 10 K€.
- Development of an information mediator. Bull, 1997, 30 K€.
- Coupling of a C++ preprocessor with an extensible schema manager. CNET-Lannion, 1993-1994 (1 year), 50 K€.
- Study of persistent data management techniques in a OO language, CNET-Lannion, 1990-1991 (1 year), 40K€.

Main national grants
- PIA Datascale on big data with Bull (leader), CEA, ActiveEon SAS, Armadillo, Twenga, IPGP, Xedix, 2013-2015, 250K€.
- PIA X-Data on big data integration with Data Publica (leader), Orange, La Poste, EDF, Cinequant, Hurence and Inria (Indes, Planete and Zenith), 125 K€.
- PREDIT EPILOG with Euxenis SAS and RISC Solutions d'Assurances, 2009-2011, 60 K€.
- ANR VERSO DataRing with Gemo (Inria Saclay), Telecom Paristech, LIG, LIRMM, 2009-2012, 200K€.
- RNTL XWiki Concerto with XPertNet, ObjectWeb, Inria, ENST, Mandriva, EISTI, 2007-2008, 150 K€.
- RNTL OpenEmbeDD with Inria, CEA-List, France Telecom R&D, AIRBUS, AnyWare, CS/SI, LAAS, Thales Avionics, Verimag, Thales Research TRT, 2006-2008, 70 K€.
- ARA Masses de Données Respire with INT, LIP6, Paris (IRISA), Regal (Inria).}{2006-2008}, 80 K€.
- ARA Masses de Données MDP2P with Gemo (Inria-Futurs), Paris (IRISA), Texmex (IRISA), 2003-2006, 80 K€.
- RNTL Legacy applications on the Net (Leg@net) Prologue Software, ASPLine, 2002-2003, 330 K€.
- Components for telecommunication network data management. MENRT, 1996-1997, 50 K€.
- Study of advanced database techniques. MRE, PRC BD3, 1991-1992, 15 K€.

Theses Advised (55)
1. Daniel Rosendo (30%), U. Rennes, 2023.
2. Renan Souza (30%), UFRJ, Rio de Janerio, 2019, **Best Ph.D. thesis award, SBBD 2021**. Now researcher at Oak Ridge National Laboratory, USA.
3. Sakina Mahboubi (40%), U. Montpellier, 2018, now Assistant Prof., University of Batna, Algeria.
4. Vitor Silva (30%), UFRJ, Rio de Janeiro, 2018, **Best Ph.D. thesis award, SBBD 2019**. Now Researcher at Apple-Siri, Los Angeles, USA.

5. Carlyna Bondiombouy (100%), U. Montpellier, 2017, now Data Scientist, Montreal, Canada.
6. Ji Liu (50%), U. Montpellier, 2016, now Technical Director, Hithink RoyalFlush Information Network Co., Hangzhou, China.
7. Valentin Leveau (30%), U. Montpellier, 2016, now postdoc at Inria.
8. Naser Ayat (40%), U. Amsterdam, 2014, now Software Architect at Connectis, Amsterdam, Netherland.
9. Miguel Liroz (50%), U. Montpellier 2, 2013, now Data Engineer at Criteo, Paris.
10. Jonas Dias (50%), UFRJ, Rio de Janeiro, 2013, **Best Ph.D. thesis award, SBBD 2015**. Now Head of Data Science at Evergen, Australia.
11. Vincenzo Gulisano (50%), U. Madrid, 2012, Now Assistant Prof., Chalmers Univ., Sweden.
12. Toufiq Sarni (40%), U. Nantes, 2011, now Research Engineer, Schneider Electrics, Nice.
13. Fady Draidi (50%), U. Montpellier 2, 2012, now Ass. Prof., Naplouse U., Palestinia.
14. Eduardo Ogasawara (50%), UFRJ, Rio de Janeiro, 2011, now Prof. CEFET, Rio de Janeiro.
15. William Kokou Dedzoe. (40%), U. Nantes, 2011, now Research Engineer, Inria Rennes.
16. Mohamed Jawad (40%), U. Nantes, 2011, now IT Consultant, Altran Technologies, Strasbourg.
17. Wenceslao Palma (50%). U. Nantes, 2010, Ass. Prof, UCV, Valparaiso, Chili.
18. Philippe Lamarre (100%), HDR, U. Nantes, 2009, now Prof. INSA Lyon.
19. Eduardo Almadeia (50%), U. Nantes, 2009, now Prof, UFP, Curritiba, Brazil.
20. Rabab Hayek (50%), U. Nantes, 2009, Gecko software, France.
21. Anthony Ventresque (50%), U. Nantes, 2009, now Prof., Trinity College Dublin, Ireland.
22. Jorge Manjarrez (50%), U. Nantes, 2009, now Associate Prof., CIC IPN, Mexico.
23. Jorge Quiané-Ruiz (50%), U. Nantes, 2008, Associate Prof., IT University of Copenhagen (ITU), deceased on May 28, 2023.
24. Reza Akbarinia (60%), U. Nantes, 2007, now Researcher, Inria. **Best Ph.D. thesis award U. Nantes, 2008.**
25. David Faye (50%), U. Nantes, 2007, now Associate Prof., U. Saint Louis, Sénégal.
26. Sandra Lemp (40%), U. Nantes, 2007, now at Acapnos, France.
27. Vidal Martins (10%), U. Nantes, 2007, now Prof., PUCPR, Curritiba, Brazil.
28. Marcos Didonet de Fabro (50%), U. Nantes, 2007, now Ass. Prof, UFP, Curritiba, Brazil.
29. Cédric Coulon (10%), U. Nantes, 2006, Vision Objects, Nantes, France.
30. Cécile Lepape (40%), UPMC, 2005, now Associate Prof. (MC) at UPMC.
31. Alexandro Bento Lima (50%), UFRJ (Rio de Janeiro), 2004, now Prof., UFRJ, Brazil.
32. Joao Pereira (20%), UVSQ, 2002, now Prof., U. Lisbon, Portugal.
33. Ioana Manolescu (40%), UVSQ, 2002, now DR Inria, head of Oak team (Inria-Saclay).
34. Helena Galhardas (20%), UVSQ, 2001, now Prof., U. Lisbon, Portugal.
35. Khaled Yagoub (50%), UVSQ, 2001, now at Oracle Corp. USA.
36. Fabio Porto (50%), PUC Rio, 2001, now Researcher, LNCC, Rio de Janeiro.
37. Olga Kapiskaia (50%), UPMC, 1999, now Prof., Audencia Business School, Nantes.
38. Hubert Naacke (50%), UVSQ, 1999, now Prof., UPMC.
39. Luc Bouganim (100%), UVSQ, 1996, now DR Inria in the PETRUS team (Inria-Saclay).
40. Daniela Florescu (100%), UPMC, 1996, **CTO and co-founder, Arico IoT, USA**.
41. Jean-Robert Gruser (40%), UPMC, 1996, now consultant, Paris.
42. Laurent Daynès (100%), UPMC, 1995, now at Sun Labs, USA.
43. Laurent Amsaleg (60%), UPMC, 1995, now DR CNRS, IRISA, Rennes.
44. Benoit Dageville (50%), UVSQ, 1995, now **co-founder, Snowflake Computing, USA**.
45. Patrick Casadessus (100%), UPMC, 1994, now at Modis, Paris.
46. Mohamed Zaït (100%), UPMC, 1994, now at Databricks., USA.
47. Eric Amiel (20%), UPMC, 1994, now software engineer, Lausanne.
48. Olivier Gruber (80%), UPMC, 1992, now Prof., U. Grenoble.
49. Marie-Jo Bellosta (100%), UPMC, 1992, now Associate Prof., U. Paris Dauphine.
50. Fabienne Viallet (80%), UPMC, 1991, now Associate Prof., U. Toulouse.
51. Mikal Ziane (80%), UPMC, 1991, now Associate Prof., U. Paris 5.

52. Björn Bergsten (80%), UPMC, 1991, now at ViaSat, Carlsbad, CA, USA.
53. Rosana Lanzelotte (80%), PUC Rio, 1991, Retired as Prof. at UNI Rio, Rio de Janeiro.
54. Philippe Pucheral (10%), UPMC, 1989, now prof., UVSQ.
55. Eric Simon (40%), UPMC, 1986, now **Technical Manager at SAP, Paris**.

Publications

Books

1. T. Özsu, P. Valduriez. Principles of Distributed Database Systems – Third Edition. Springer, 850 pages, 2011. Tsinghua University Press (Chinese Edition), 2013. Fourth Edition, Springer, 674 pages, 2020.
2. T. Özsu, P. Valduriez. Principles of Distributed Database Systems. Prentice Hall, Englewood Cliffs, New Jersey, 1st edition, 562 pages, 1991; 2nd edition, 666 pages, 1999.
3. M. Bouzeghoub, G. Gardarin, P. Valduriez. Object Technology: concepts and methods. International Thomson Computer Press, 382 pages, 1997.
4. M. Bouzeghoub, G. Gardarin, P. Valduriez. Les Objets. Eyrolles, Paris, 1st edition, 344 pages, 1994; 2nd edition, 450 pages, 1997.
5. G. Gardarin, P. Valduriez. Relational Databases and Knowledge Bases. Addison-Wesley Publishing Company, Reading, Massachusetts, 448 pages, 1990.
6. P. Valduriez, G. Gardarin. Analysis and Comparison of Relational Database Systems. Addison-Wesley Publishing Company, Reading, Massachusetts, 280 pages, 1990.
7. T. Özsu, U. Dayal, P. Valduriez (eds.). Distributed Object Management. Morgan Kaufmann, 441 pages, 1993.
8. P. Valduriez (ed.). Data Management and Parallel Processing. Chapman and Hall, London, 350 pages, 1992.

Book chapters

9. T. Morzy, P. Valduriez, L. Bellatreche. Advances in Databases and Information Systems. 19th East European Conference, ADBIS 2015, Poitiers, France, September 8-11, 2015, LNCS 9282, Springer 2015.
10. T. Özsu, P. Valduriez. Distributed and Parallel Database Systems. In Information Systems and Information Technology, Volume 2 (Computing Handbook, Third Edition) H. Topi and A. Tucker (eds). Chapman and Hall, 2014.
11. P. Valduriez. Parallel Database Management. Encyclopedia of Database Systems. L. Liu and T. Özsu (eds), Springer, 2026-2029, 2009.
12. P. Valduriez. Parallel Data Placement. Encyclopedia of Database Systems. L. Liu and T. Özsu (eds), Springer, 2024-2026, 2009.
13. P. Valduriez, E. Pacitti. Parallel Database Systems. Handbook of Database Technology, J. Hammer and M. Scheider (eds.), CRC Press, 2007.
14. R. Akbarinia, V. Martins, E. Pacitti, P. Valduriez. Design and Implementation of APPA. Global Data Management, R. Baldoni, G. Cortese, F. Davide (eds.), IOS Press, 2006.
15. T. Özsu, P. Valduriez. Distributed and Parallel Database Systems. Handbook of Computer Science and Engineering, 2nd Edition, A. Tucker (ed.), CRC Press, 2004.
16. P. Valduriez, M. Zaït, M. Ziane . Designing an Optimizer for Parallel Relational Systems. Parallel Database Techniques, M. Abdelguerfi and K-F. Wong (eds.), IEEE Computer Society Press, 1998.
17. L. Daynès, M. Atkinson, P. Valduriez . Customizable Concurrency Control for Persistent Java. Advanced Transaction Models and Architectures, S. Jajodia and L. Kerschberg (eds.), Kluwer Academic Publishers, 1997.
18. T. Özsu, P. Valduriez. Distributed and Parallel Database Systems. Handbook of Computer Science and Engineering, A. Tucker (ed.), CRC Press, 1996.

19. T. Özsu, P. Valduriez. Distributed Data Management: unsolved problems and new issues. Readings in Distributed Computing, T. Casavant and M. Singhal (eds.), IEEE Computer Society Press, 1993.
20. T. Özsu, U. Dayal, P. Valduriez. An Introduction to Distributed Object Management. Distributed Object Management,T. Özsu, U. Dayal and P. Valduriez (eds.), Morgan-Kaufmann, 1993.
21. O. Gruber, P. Valduriez. Object Management in Shared-Memory Parallel Data Servers. Data Management and Parallel Processing, P. Valduriez (ed.), Chapman and Hall, London, 1992.
22. S. Khoshafian, P. Valduriez. Persistence, Sharing and Object Orientation. a database perspective. Advances in Database Programming Languages, F. Bancilhon and P. Buneman (eds.), ACM Press, 1990.
23. S. Khoshafian, P. Valduriez. Parallel Execution Strategies for Declustered Databases. Database Machines and Knowledge Base Machines, M. Kitsuregawa and H. Tanaka (eds.), Kluwer Academic Publishers, 1988.
24. G. Gardarin, P. Bernadat, P. Valduriez, Y. Viemont. SABRE. A Database System for a Multi-Microprocessor Database Machine. Advanced Database Machine Architecture, D. Hsiao (ed.), Prentice Hall, 1983.
25. P. Valduriez, G. Gardarin. Multiprocessor Join Algorithms of Relations. Improving Database Usability and Responsiveness, P. Scheuermann (ed.), Academic Press, 1982.

Papers published in international journals

1. R. Zorrilla, E. Ogasawara, P. Valduriez, F. Porto. A Data-Driven Model Selection Approach to Spatio-Temporal Prediction. Trans. Large Scale Data Knowledge-Centered Systems (TLDKS), Volume 14790, 98-118, 2024.
2. R. Saldanha, R. Akbarinia, M. Pedroso, V. Ribeiro, C. Cardoso, E. Pena, P. Valduriez, F. Porto. Zonal Statistics Datasets of Climate Indicators for Brazilian Municipalities. Environmental Data Science, 2024, 3, 2024.
3. R. Akbarinia, C. Botella, A. Joly, F. Masseglia, M. Mattoso, E. Ogasawara, D. de Oliveira, E. Pacitti, F. Porto, C. Pradal, D. Shasha, P. Valduriez: Life Science Workflow Services (LifeSWS): motivations and architecture. Trans. on Large-Scale Data- and Knowledge-Centered Systems (TLDKS), 55: 1-24, 2023.
4. J. Liu, D. Dong, X. Wang, A. Qin, X. Li, P. Valduriez, D. Dou, D. Yu. Large-scale Knowledge Distillation with Elastic Heterogeneous Computing Resources. Concurrency and Computation: Practice and Experience, 35(26) .e7272, 2023.
5. D. Rosendo, A. Costan, P. Valduriez, G. Antoniu. Distributed intelligence on the Edge-to-Cloud Continuum: A systematic literature review. Journal of Parallel and Distributed Computing, Elsevier, 2022, 166, pp.71-94.
6. J. Liu, C. Bondiombouy, L. Mo, P. Valduriez. Two-Phase Scheduling for Efficient Vehicle Sharing. IEEE Transactions on Intelligent Transportation Systems, 2022, 23 (1), pp.457-470.
7. R. Jimenez-Peris, D. Burgos-Sancho, F. Ballesteros, M. Patiño-Martinez, P. Valduriez. Elastic scalable transaction processing in LeanXcale. Information Systems, Elsevier, 2022, pp. 1-18.
8. R. Souza, V. Silva, A. Lima, D. de Oliveira, P. Valduriez, M. Mattoso: Distributed In-memory Data Management for Workflow Executions, PeerJ Computer Science, 2021.
9. G. Heidsieck, D. de Oliveira, E. Pacitti, C. Pradal, F. Tardieu, P. Valduriez. Cache-aware Scheduling of Scientific Workflows in Multisite Cloud. Future Generation Computer Systems, 2021.
10. B. Kolev, O. Levchenko, E. Pacitti, P. Valduriez, R. Jiménez-Peris, P. Kranas, M. Patiño-Martinez. Parallel Query Processing in a Polystore. Distributed and Parallel Databases, 36 pages, 2021.
11. P. Valduriez, R. Jimenez-Peris, T. Özsu. Distributed Database Systems: The Case for NewSQL. Trans. on Large-Scale Data- and Knowledge-Centered Systems (TLDKS), 15 pages, 2021.
12. V. Silva, V. Campos, T. Guedes, J. Camata, D. de Oliveira, A. Coutinho, P. Valduriez, M. Mattoso. DfAnalyzer: Runtime dataflow analysis tool for Computational Science and Engineering applications. SoftwareX, Elsevier, 12 pages, 2020.

13. G. Heidsieck, D.de Oliveira, E. Pacitti, C. Pradal, F. Tardieu, P. Valduriez. Efficient Execution of Scientific Workflows in the Cloud Through Adaptive Caching. Transactions on Large-Scale Data- and Knowledge-Centered Systems, 44: 41-66, 2020.
14. N. Lemus, F. Porto, Y. Souto, R. Pereira, J. Liu, E. Pacitti, P. Valduriez. SUQ2: Uncertainty Quantification Queries over Large Spatio-temporal Simulations. Bulletin of the Technical Committee on Data Engineering, IEEE Computer Society, 43(1):47-59, 2020.
15. J. Liu, N. Lemus, E. Pacitti, F. Porto, P. Valduriez. Parallel Computation of PDFs on Big Spatial Data Using Spark. Distributed and Parallel Databases, 38, 63-100, 2020.
16. R. Souza, V. Silva, A. Coutinho, P. Valduriez, M. Mattoso. Data Reduction in Scientific Workflows using Provenance Monitoring and User Steering. Future Generation Computer Systems (FGCS), Elsevier, 110, 481-501, 2020.
17. R. Souza, V. Silva, J. Camata, A. Coutinho, P. Valduriez, M. Mattoso. Keeping Track of User Steering Actions in Dynamic Workflows. Future Generation Computer Systems (FGCS), Elsevier, 99, pp.624-643, 2019.
18. J. Liu, L. Pineda, E. Pacitti, A. Costan, P. Valduriez, G. Antoniu, M. Mattoso. Efficient Scheduling of Scientific Workflows using Hot Metadata in a Multisite Cloud. IEEE Trans. on Knowledge and Data Engineering (TKDE), 31(10): 1940-1953, 2018.
19. V. Silva, D. de Oliveira, P. Valduriez, M. Mattoso. DfAnalyzer: Runtime Dataflow Analysis of Scientific Applications using Provenance. Proceedings of the VLDB Endowment (PVLDB), 11(12): 2082-2085, 2018.
20. D. Yagoubi, R. Akbarinia, B. Kolev, O. Levchenko, F. Masseglia, P. Valduriez, D. Shasha. ParCorr: efficient parallel methods to identify similar time series pairs across sliding windows. Data Mining and Knowledge Discovery 32(5): 1481-1507, 2018.
21. J. Liu, E. Pacitti, P. Valduriez. A Survey of Scheduling Frameworks in Big Data Systems. Int. Journal of Cloud Computing, 7(2): 103-128, 2018.
22. J. Camata, V. Sousa, P. Valduriez, M. Mattoso, A. Coutinho. In situ visualization and data analysis for turbidity currents simulation. Computers & Geosciences 110: 23-31, 2018.
23. C. Pradal, S. Artzet, J. Chopard, D. Dupuis, C. Fournier, M. Mielewczik, V. Nègre, P. Neveu, D. Parigot, P. Valduriez, S. Cohen Boulakia. InfraPhenoGrid: A scientific workflow infrastructure for plant phenomics on the Grid. Future Generation Computer Systems, Elsevier, 67: 341-353, 2017.
24. V. Sousa, J. Leite, J. Camata, D. de Oliveira, A. Coutinho, P. Valduriez, M. Mattoso. Raw data queries during data-intensive parallel workflow execution. Future Generation Computer Systems, Elsevier, 75: 402-422, 2017.
25. J. Liu, E. Pacitti, P. Valduriez, M. Mattoso. Scientific Workflow Scheduling with Provenance Data in a Multisite Cloud. Trans. on Large-Scale Data- and Knowledge-Centered Systems (TLDKS), 33: 80-112, 2017.
26. C. Bondiombouy, B. Kolev, O. Levchenko, P. Valduriez. Multistore Big Data Integration with CloudMdsQL. Trans on Large-Scale Data and Knowledge-Centered Systems, Springer, 9940:48-74, 2016.
27. C. Bondiombouy, P. Valduriez. Query Processing in Multistore Systems: an overview. Int. Journal of Cloud Computing, 5(4): 309-346, 2016.
28. H. Lustosa, F. Porto, P. Blanco, P. Valduriez. Database System Support of Simulation Data. Proceedings of the VLDB Endowment (PVLDB), 9(13): 1329-1340, 2016.
29. J. Liu, E. Pacitti, P. Valduriez, M. Mattoso. Multi-Objective Scheduling of Scientific Workflows in Multisite Clouds. Future Generation Computer Systems, Elsevier, 63: 76-95, 2016.
30. B. Kolev, P. Valduriez, C. Bondiombouy, R. Jiménez-Peris, R. Pau, J. Pereira. CloudMdsQL: Querying Heterogeneous Cloud Data Stores with a Common Language. Distributed and Parallel Databases, 34(4): 463-503, 2016.
31. V. Silva, D. de Oliveira, P. Valduriez, M. Mattoso. Analyzing Related Raw Data Files through Dataflows. Concurrency and Computation: Practice and Experience, 28(8): 2528-2545, 2016.
32. M. Liroz-Gistau, R. Akbarinia, D. Agrawal, P. Valduriez. FP-Hadoop: Efficient processing of skewed MapReduce jobs. Information Systems, 60: 69-84, 2016.

33. M. Liroz-Gistau, R. Akbarinia, P. Valduriez. FP-Hadoop: Efficient Execution of Parallel Jobs Over Skewed Data. Proceedings of the VLDB Endowment (PVLDB), 8(12): 1856-1867, 2015.
34. J. Liu, E. Pacitti, P. Valduriez, M. Mattoso. A Survey of Data-Intensive Scientific Workflow Management. Journal of Grid Computing, 13(4):457-493, 2015.
35. J. Dias, G. Guerra, F. Rochinha, A. Coutinho, P. Valduriez, M. Mattoso. Data-Centric Iteration in Dynamic Workflows. Future Generation Computer Systems, Elsevier, Vol. 4, 114-126, 2015.
36. N. Ayat, R. Akbarinia, H. Afsarmanesh, P. Valduriez. Entity resolution for probabilistic data. Information Sciences 277: 492-511, 2014.
37. Aboulnaga, B. Chin Ooi, P. Valduriez. Special section on data-intensive cloud infrastructure. VLDB J. 23(6): 843, 2014.
38. M. Liroz-Gistau, R. Akbarinia, E. Pacitti, F. Porto, P. Valduriez. Dynamic Worlkload-based Partitioning Algorithms for Continuously Growing Databases. Trans on Large-Scale Data and Knowledge-Centered Systems, Springer, 12:105-128, 2013.
39. J. Meira, E. Cunha de Almeida, G. Sunyé, Y. Le Traon, P. Valduriez. Stress Testing of Transactional Database Systems. Journal of Information and Data Management, 4(3): 279-294, 2013.
40. E. Ogasawara, J. Dias, V. Silva, F. Chirigati, D. de Oliveira, F. Porto, P. Valduriez, M. Mattoso. Chiron: A Parallel Engine for Algebraic Scientific Workflows. Concurrency and Computation: Practice and Experience, 25(16):2327-2341, 2013.
41. N. Ayat, R. Akbarinia , H. Afsarmanesh, P. Valduriez. Entity Resolution for Distributed Probabilistic Data. Distributed and Parallel Databases, 31(4):509-542, 2013.
42. W. Kokou Dedzoe, P. Lamarre, R. Akbarinia, P. Valduriez. As-Soon-As-Possible Top-k Query Processing in P2P Systems. Trans on Large-Scale Data and Knowledge-Centered Systems, 9:1-27, 2013.
43. K. Park, P. Valduriez. A Hierarchical Grid Index (HGI), Spatial Queries in Wireless Data Broadcasting. Distributed and Parallel Databases, 31(3): 413-446, 2013.
44. R. Akbarinia, P. Valduriez, G. Verger. Efficient Evaluation of SUM Queries Over Probabilistic Data. IEEE Trans. on Knowledge and Data Engineering, 25(4): 764-775, 2013.
45. V. Gulisano, R. Jimenez-Peris, M. Patiño-Martinez, C. Soriente, P. Valduriez. StreamCloud: An Elastic and Scalable Data Streaming System. IEEE Trans. On Parallel and Distributed Systems, 23(12): 2351 – 2365, 2012.
46. V. Gulisano, R. Jimenez-Peris, M. Patiño-Martinez, C. Soriente, P. Valduriez. A Big Data Platform for Large Scale Event Processing. ERCIM News (89), 2012.
47. E. Pacitti, P. Valduriez. Zenith: Scientific Data Management on a Large Scale. ERCIM News (89), 2012.
48. J. Quiané-Ruiz, P. Lamarre, P. Valduriez. Satisfaction-based Query Replication: an automatic and self-adaptable approach for replicating queries in the presence of autonomous participants. Distributed and Parallel Databases, 30(1):1-26, 2012.
49. E. Ogasawara, J. Dias, D. Oliveira, F. Porto, P. Valduriez, M. Mattoso. An Algebraic Approach for Data-centric Scientific Workflows. Proceedings of the VLDB Endowment (PVLDB), 4(12):1328-1339, 2011.
50. K. Park, P. Valduriez. Energy Efficient Data Access in Mobile P2P Networks. IEEE Trans. on Knowledge and Data Engineering, 23(11):1619-1634, 2011.
51. R. Akbarinia, E. Pacitti, P. Valduriez. Best Position Algorithms for Efficient Top-k Query Processing. Information Systems, 36(6):973-989, 2011.
52. R. Akbarinia, M. Tlili, E. Pacitti, P. Valduriez, A. A.B. Lima. Continuous Timestamping for Efficient Replication Management in DHTs. Trans. on Large-Scale Data and Knowledge-Centered Systems, Springer, 2011.
53. G. A. Vouros, A. Papasalouros, K. Tzonas, A. Valarakos, K. Kotis, P. Lamarre, J. Quiané-Ruiz, P. Valduriez. A Semantic Information System for Services and Traded Resources in Grid e-Markets. Future Generation Computer Systems, 26(7): 916-933, 2010.

54. E. Cunha de Almeida, G. Sunye, Y. Le Traon, P. Valduriez. Testing Peer-to-peer Systems. Empirical Software Engineering (Springer), 15(4):346-379, 2010, invited as extended version of ISSRE2008 (second best paper).
55. A. Lima, M. Mattoso, P. Valduriez. Adaptive Virtual Partitioning for OLAP Query Processing in a Database Cluster. Journal of Information and Data Management (JIDM), 1(1): 75-88, 2010.
56. A. Lima, M. Mattoso, P. Valduriez. Adaptive Virtual Partitioning: Further Developments. Journal of Information and Data Management (JIDM), 1(1): 89-92, 2010.
57. K. Park, H. Choo, P. Valduriez. A Scalable Energy-efficient Continuous Nearest Neighbor Search in Wireless Broadcast Systems. Wireless Networks, 16(4): 1011-1031, 2010.
58. W. Palma, R. Akbarinia, E. Pacitti, P. Valduriez. DHTJoin: processing continuous join queries using DHT networks. Distributed and Parallel Databases, 26(2-3): 291-317, 2009.
59. J. Quiané-Ruiz, P. Lamarre, P. Valduriez. A Self-Adaptable Query Allocation Framework for Distributed Information Systems. The VLDB Journal, 18(3):649-674, 2009.
60. M. Didonet Del Fabro, Patrick Valduriez. Towards the Efficient Development of Model Transformations using Model Weaving and Matching Transformations. Software and Systems Modeling, 8(3): 305-324, 2009.
61. A. Lima, C. Furtado, P. Valduriez, M. Mattoso. Parallel OLAP Query Processing in Database Clusters with Data Replication. Distributed and Parallel Databases, 25(1-2): 97-123, 2009.
62. M. Tlili, W. K. Dedzoe, E. Pacitti, R. Akbarinia P. Valduriez, et al. P2P Logging and Timestampg for Reconciliation, Proceedings of the VLDB Endowment (PVLDB), 1(2): 1420-1423, 2008.
63. N. Anciaux, L. Bouganim, P. Pucheral, P. Valduriez. DiSC: Benchmarking Secure Chip DBMS. IEEE Trans. on Knowledge and Data Engineering, 20(10):1363-1377, 2008.
64. N. Kotowski, A. Lima, E. Pacitti, P. Valduriez, M. Mattoso. OLAP Query Processing in Grids, Concurrency and Computation: Practice and Experience, Special Issue on Best Papers of the VLDB Data Management in Grids Workshop (VLDB DMG 2007), 20(17):2039-2048, 2008.
65. R. Hayek, G. Raschia, P. Valduriez, N. Mouaddib.Towards a Data Summarization Service in APPA. Int. Journal of Pervasive Computing and Communications (IJPCC), 4(4):390-410, 2008.
66. A. Furtado, A. Lima, E. Pacitti, P. Valduriez, M. Mattoso. Adaptive Hybrid Partitioning for OLAP Query Processing in a Database Cluster. Int. Journal of High Performance Computing and Networking, Special issue on Best Papers from SBAC2005, 5(4):251-262, 2008.
67. P. Lamarre, S. Lemp, S. Cazalens, P. Valduriez. A Flexible Mediation Process for Large Distributed Information Systems. Int. Journal of Cooperative Information Systems, 16(2):299-332, 2007.
68. E. Pacitti, P. Valduriez, M. Mattoso. Grid Data Management: open problems and new issues. Journal of Grid Computing, 5(3):273-281, 2007.
69. S. Gançarski, H. Naacke, E. Pacitti, P. Valduriez. The Leganet System. Freshness-Aware Transaction Routing in a Database Cluster. Information Systems, (32)2:320-343, 2007.
70. R. Akbarinia, E. Pacitti, P. Valduriez. Reducing Network Traffic in Unstructured P2P Systems Using Top-K Queries. Distributed and Parallel Databases, 19(2):67-86, 2006.
71. E. Pacitti, C. Coulon, P. Valduriez, T. Özsu. Preventive Replication in a Database Cluster. Distributed and Parallel Databases, 18(3):223-251, 2005.
72. C. Lepape, S. Gançarski, P. Valduriez. Data Quality in a Database Cluster with Lazy Replication. Journal of Distributed Information Management, Special Issue on Distributed Data Management, 3(2):82-87, 2005.
73. G. Bernard, J. Ben-Othman, L. Bouganim, G. Canals, B. Defude, J. Ferrié, S.Gançarski, R. Guerraoui, P. Molli, P. Pucheral, C. Roncancio, P. Serrano-Alvarado, P. Valduriez. Mobile Databases. a report on open issues and research directions. ACM SIGMOD Record, 33(2):78-83, 2004.
74. P. Pucheral, L. Bouganim, P. Valduriez, C. Bobineau. PicoDBMS: Scaling down database techniques for the Smartcard. The VLDB Journal, Special issue on Best Papers from VLDB2000, 10(2-3):120-132, 2001.
75. A. Tanaka, P. Valduriez, et al. The Ecobase Environmental Information System. applications, architecture and open issues. ACM SIGMOD Record, 30(3):70-75, 2001.

76. A. Tanaka, P. Valduriez, et al.. The Ecobase Project. Database and Web Technologies for Environmental Information Systems. Network and Information Systems Journal, 3(5-6):865-882, 2000.
77. L. Bouganim, F. Fabret, C. Mohan, P. Valduriez. A Dynamic Query Processing Architecture for Data Integration Systems. IEEE Data Engineering Bulletin, 23(2):42-48, 2000.
78. H. Naacke, A. Tomasic, P. Valduriez. Validating Mediator Cost Models with DISCO. Network and Information Systems Journal, Hermès, 2(5-6):639-663, 1999.
79. L. Bouganim, D. Florescu, P. Valduriez . Load Balancing for Parallel Query Execution on NUMA Multiprocessors. Distributed and Parallel Databases, 7(1):99-121, 1999.
80. L. Bouganim, O. Kapiskaia, P. Valduriez. Dynamic Memory Allocation for Large Query Execution. Networking and Information Systems Journal, Special issue on the best papers from CIKM'98, 1(6):629-652, 1998.
81. A. Tomasic, L. Raschid, P. Valduriez. Scaling Access to Heterogeneous Data Sources with DISCO. IEEE Trans. on Knowledge and Data Engineering, Extended version of best papers from ICDCS'96, 10(5):808-823, 1998.
82. M. Skubiszewski, P. Valduriez. Using GC-Consistent Cuts for Concurrent Garbage Collection in O2. Networking and Information Systems, 1(2-3):213-230 (1998)
83. E. Pacitti, P. Valduriez. Replicated Databases: concepts, architectures and techniques. Networking and Information Systems, 1(3):519-546, 1998.
84. T. Özsu, P. Valduriez. Distributed and Parallel Database Systems - Technology and current state-of-the-art. ACM Computing Surveys, 28(1):125-128, 1996.
85. W. Hasan, D. Florescu, P. Valduriez. Open Issues in Parallel Query Optimization. ACM SIGMOD Record, 25(3):28-33, 1996.
86. M-J. Bellosta, G. Lecorgne, B. Nicolas, P. Valduriez. Extending an ODBMS to Support ISO Network Management. Int. Journal of Intelligent Systems for Electrical Engineering and Communication. July 1996.
87. D. Shasha, F. Llirbat, E. Simon, P. Valduriez. Transaction Chopping: Algorithms and Performance Studies. ACM Trans. on Database Systems, 20(3):325-363, 1995.
88. D. Florescu, J-R. Gruser, M. Novak, P. Valduriez, M. Ziane. Design and Implementation of Flora, a Language for Object Algebra. Information Science, 87(1-3):1-27,1995.
89. D. Florescu, P. Valduriez. Rule-based Query Rewriting in the Flora Optimizer. IEICE Trans. on Information and Systems, 8, 1995.
90. M. Zait, P. Valduriez, D. Florescu. Validating a Parallel Query Optimizer. Engineering of Information Systems, (3)1, 1995.
91. R. Lanzelotte, P. Valduriez, M. Zait, M. Ziane. Industrial-strenght Parallel Query Optimization: issues and lessons. Information Systems, (19)4:311-330, 1994.
92. B. Bergsten, M. Couprie, P. Valduriez. Overview of Parallel Architectures for Databases, The Computer Journal, (36)8:734-740, 1993.
93. P. Valduriez. Parallel Database Systems. open problems and new issues. Distributed and Parallel Databases, 1(2):137-165, 1993.
94. S. Danforth, P. Valduriez. A FAD for Data-Intensive Applications. IEEE Trans. on Data and Knowledge Engineering, 4(1):34-51, 1992.
95. P. Valduriez, S. Danforth. Functional SQL, an SQL Upward Compatible Database Programming Language. Information Sciences, 62(3):183-203, 1992.
96. S. Danforth, P. Valduriez. The Data Model of FAD, a Database Programming Language. Information Sciences, (60)1&2 :51-75, 1992.
97. T. Özsu, P. Valduriez. Systèmes de Gestion de Bases de Données Répartis et Paralléles. Canadian Information Processing/Informatique Canadienne, 1992.
98. T. Özsu, P. Valduriez. Distributed Database Systems. Where Are We Now? IEEE Computer, (24)8:68-78, 1991.
99. H. Boral, P. Valduriez, et al.. Prototyping Bubba, a Highly Parallel Database System. IEEE Trans. on Knowledge and Data Engineering, (2)1:4-24, 1990.

100. B. Hart, P. Valduriez, S. Danforth. Parallelizing FAD using Compile Time Analysis Techniques. IEEE Bulletin on Database Engineering, (12)1:9-15, 1989.
101. P. Valduriez, S. Khoshafian. Parallel Evaluation of the Transitive Closure of a Database Relation. Int. Journal of Parallel Programming, (17)1:19-37, 1988.
102. P. Valduriez. Complex Objects in Relational Database Systems. Technology and Science of Informatics, (6)5:597-608, 1987.
103. P. Valduriez. Join Indices. ACM Trans. on Database Systems, (12)2:218-246, 1987.
104. P. Valduriez. Optimization of Complex Queries Using Join Indices. IEEE Bulletin on Database Engineering, (9)4:10-16, 1986.
105. Gardarin, P. Valduriez, et al.. SABRINA, a Relational Database System Developed in a Research Environment. Technology and Science of Informatics, (5)6:453-474, 1986.
106. P. Valduriez, G. Gardarin. Join and Semi-join Algorithms for a Multiprocessor Database Machine. ACM Trans. on Database Systems, (9)1:133-161, 1984.

Papers published in international conferences

1. J. Liu, J. Ren, R. Jin, Z. Zhang, Y. Zhou, P. Valduriez, D. Dou. Fisher Information-based Efficient Curriculum Federated Learning with Large Language Models. Conference on Empirical Methods in Natural Language Processing (EMNLP), 10497-10523, 2024.
2. R. Saldanha, V. Ribeiro, E. Pena, M. Pedroso, R. Akbarinia, P. Valduriez, F. Porto. Subset Models for Multivariate Time Series Forecast. Int. Conf. on Data Engineering Workshops (ICDEW), 86-90, 2024.
3. J. Liu, T. Che, Y. Zhou, R. Jin, H. Dai, D. Dou, P. Valduriez. AEDFL: Efficient Asynchronous Decentralized Federated Learning with Heterogeneous Devices. SIAM Int. Conf. on Data Mining (SDM) 833-841, 2024.Rosendo, M. Mattoso, A. Costan, R. Souza, D. Pina, P. Valduriez, G. Antoniu. ProvLight: Efficient Workflow Provenance Capture on the Edge-to-Cloud Continuum. IEEE International Conference on Cluster Computing (Cluster), 2023.
4. D. Rosendo, K. Keahey, A. Costan, M. Simonin, P. Valduriez, G. Antoniu. KheOps: Cost-effective Repeatability, Reproducibility, and Replicability of Edge-to-Cloud Experiments.
5. ACM Conference on Reproducibility and Replicability (REP), 2023.
6. V. Ribeiro, E. Pena, R. Saldanha, R. Akbarinia, P. Valduriez, F. Arif, J. Stoyanovich, F. Porto. Subset Modelling: A Domain Partitioning Strategy for Data-efficient Machine-Learning
7. Simpósio Brasileiro de Banco de Dados (SBBD), 2023.
8. P. Valduriez, F. Porto. Data and Machine Learning Model Management with Gypscie. CARLA Workshop on HPC and Data Sciences meet Scientific Computing, SCALAC, 2022.
9. C. Silva, P. Valduriez, F. Porto. Integrating Machine Learning Model Ensembles to the SAVIME Database System. Brazilian Symposium on Databases (SBBD), 2022.
10. R. Zorrilla, E. Ogasawara, P. Valduriez, F. Porto. A Data-Driven Model Selection Approach to Spatio-Temporal Prediction. Brazilian Symposium on Databases (SBBD), 2022.
11. D. Rosendo, A. Costan, G. Antoniu, M. Simonin, J-C. Lombardo, A. Joly, P. Valduriez. Reproducible Performance Optimization of Complex Applications on the Edge-to-Cloud Continuum. IEEE International Conference on Cluster Computing (Cluster), 2021.
12. G. Heidsieck, D. de Oliveira, E. Pacitti, C. Pradal, F. Tardieu, P. Valduriez. Distributed Caching of Scientific Workflows in Multisite Cloud. International Conference on Database and Expert Systems Applications (DEXA), Best paper award, 2020.
13. G. Heidsieck, D. de Oliveira, E. Pacitti, C. Pradal, F. Tardieu, P. Valduriez. Adaptive Caching for Data-Intensive Scientific Workflows in the Cloud. International Conference on Database and Expert Systems Applications (DEXA), 2019.
14. H. Lustosa, F. Porto, P. Valduriez. SAVIME: A Database Management System for Simulation Data Analysis and Visualization. Brazilian Symposium on Databases (SBBD), 2019.
15. O. Levchenko, B. Kolev, D-E. Yagoubi, D. Shasha, T. Palpanas, P. Valduriez, R. Akbarinia, F. Masseglia. Distributed Algorithms to Find Similar Time Series. European Conference on Machine Learning and Knowledge Discovery in Databases (ECML-PKDD), pp.781-785, 2019.

16. R. Souza, L. Azevedo, R. Thiago, E. Soares, M. Nery, M. Netto, E. Vital Brazil, R. Cerqueira, P. Valduriez, M. Mattoso.Efficient Runtime Capture of Multiworkflow Data Using Provenance. International eScience Conference, pp.1-10, 2019.
17. H. Lustosa, F. Porto, P. Valduriez: SAVIME: A Database Management System for Simulation Data Analysis and Visualization. Brazilian Symposium on Databases (SBBD), pp.1-12, 2019.
18. B. Kolev, O. Levchenko, E. Pacitti, P. Valduriez, R. Vilaça, R. Gonçalves, R. Jiménez-Peris, P. Kranas. Parallel Polyglot Query Processing on Heterogeneous Cloud Data Stores with LeanXcale. IEEE Big Data Conference, 1757-1766, 2019.
19. F. Porto, J. Rittmeyer, E. Ogasawara, A. Krone-Martins, P. Valduriez, D. Shasha. Point Pattern Search in Big Data. Int. Conf. on Scientific and Statistical Database Management (SSDBMS), 21:1-21:12, 2018.
20. S. Mahboubi, R. Akbarinia, P. Valduriez. Privacy-Preserving Top-k Query Processing in Distributed Systems. European Conf. on Parallel and Distributed Computing (Euro-Par), 281-292, 2018.
21. S. Mahboubi, R. Akbarinia, P. Valduriez. Answering Top-k Queries over Outsourced Sensitive Data in the Cloud. Int. Conf. on Databases and Expert Systems Applications (DEXA), 218-231, 2018.
22. S. Mahboubi, R. Akbarinia, P. Valduriez. Top-k Query Processing over Distributed Sensitive Data. Int. Database Engineering & Applications Symposium (IDEAS), 208-216, 2018.
23. J. Liu, N. Lemus, E. Pacitti, F. Porto, P. Valduriez. Computation of PDFs on Big Spatial Data: Problem & Architecture. Latin America Data Science Workshop (LADaS), in conjunction with VLDB2018, 6 pages, 2018.
24. P. Valduriez, et al. Scientific Data Analysis Using Data-Intensive Scalable Computing: the SciDISC Project. Latin America Data Science Workshop (LADaS), in conjunction with VLDB2018, 8 pages, 2018.
25. F. Porto, A. Khatibi, J. Rittmeyer, E. Ogasawara, P. Valduriez, D. Shasha. Constellation Queries over Big Data. Brazilian Symposium on Databases (SBBD), 85-96, 2018.
26. R. Souza, V. Silva, P. Miranda, A. Lima, P. Valduriez, M. Mattoso. Spark Scalability Analysis in a Scientific Workflow. Brazilian Symposium on Databases (SBBD), Best paper award, 288-293, 2017.
27. A. Khatibi, F. Porto, J. Rittmeyer, E. Ogasawara, P. Valduriez, D. Shasha. Pre-processing and Indexing Techniques for Constellation Queries in Big Data. Int. Conf. on Big Data Analytics and Knowledge Discovery (DaWaK), 164-172, 2017.
28. H. Lustosa, N. Lemus, F. Porto, P. Valduriez. RS: An Array Model with Rich Semantics for Multidimensional Data. ER FORUM 2017: Conceptual Modeling : Research In Progress, 114-127, 2017.
29. J. Camata, P. Valduriez et al. Enhancing Energy Production with Exascale HPC Methods. High Performance Computing - Third Latin American Conference (CARLA), Springer, Communications in Computer and Information Science 697, 233-246, 2017.
30. L. Pineda-Morales, J. Liu , A. Costany, E Pacitti, G. Antoniu, P. Valduriez, M. Mattoso. Managing Hot Metadata for Scientific Workflows on Multisite Clouds. IEEE Big Data Conference, 390-397, 2016.
31. B. Kolev, C. Bondiombouy, P. Valduriez, R. Jiménez-Peris, R. Pau, J. Pereira. The CloudMdsQL Multistore System. ACM SIGMOD Int. Conf. on Management of Data, 2113-2116, 2016.
32. B. Kolev, R. Pau, O. Levchenko, P. Valduriez, R. Jiménez-Peris, J. Pereira. Benchmarking polystores: The CloudMdsQL experience. Workshop on Methods to Manage Heterogeneous Big Data and Polystore Databases, IEEE Big Data Conference, 2574-2579, 2016.
33. B. Kolev, C. Bondiombouy, O. Levchenko, P. Valduriez, R. Jiménez-Peris, R. Pau, J. Pereira. Design and Implementation of the CloudMdsQL Multistore System. Int. Conf. on Cloud Computing and Services Science (CLOSER), 352-359, 2016.
34. J. Liu, E. Pacitti, P. Valduriez, M. Mattoso. Scientific Workflow Scheduling with Provenance Support in Multisite Cloud. Int. Conf. on High Performance Computing for Computational Science (VecPar), 2016.

35. R. García, J. Camata, J. Cela, D. Costa, A. Coutinho, D. Fernández-Galisteo, C. Jiménez, V. Kourdioumov, M. Mattoso, T. Miras, J. Moríñigo, J. Navarro, P. Navaux, D. de Oliveira, M. Pascual, V. Sousa, R. Souza, P. Valduriez. Enhancing Energy Production with Exascale HPC Methods. High Performance Computing - Third Latin American Conference, CARLA 2016, Mexico City, Mexico, Revised Selected Papers. Communications in Computer and Information Science 697, 233-246, 2016.
36. V. Leveau, A. Joly, O. Buisson, P. Valduriez. Spatially Localized Visual Dictionary Learning. Int. Conf. on Multimedia Retrieval (ICMR), 367-370, 2016.
37. R. Souza, V. Silva, D. de Oliveira, P. Valduriez, A. Lima, M. Mattoso. Parallel Execution of Workflows Driven by a Distributed Database Management System. Int. Conf. For High Performance Computing, Networking, Storage and Analysis (SC15), 2015.
38. H. Lustosa, F. Porto, R. Costa, P. Blanco, P. Valduriez. Managing Simulation Data with Multidimensional Arrays. Brazilian Symposium on Databases (SBBD), 2015.
39. C. Bondiombouy, B. Kolev, O. Levchenko, P. Valduriez. Integrating Big Data and Relational Data with a Functional SQL-like Query Language. Int. Conf. on Databases and Expert Systems Applications (DEXA), 2015.
40. R. Akbarinia, M. Liroz-Gistau, D. Agrawal, P. Valduriez. An Efficient Solution for Processing Skewed MapReduce Jobs. Int. Conf. on Data Management in Grid and P2P Systems (Globe), 417-429, 2015.
41. C. Pradal, C. Fournier, P. Valduriez, S. Cohen-Boulakia. OpenAlea: scientific workflows combining data analysis and simulation. Int. Conf. on Scientific and Statistical Database Management (SSDBMS), 11:1-11:6, 2015.
42. V. Leveau, A. Joly, O. Buisson, P. Valduriez. Kernelizing Spatially Consistent Visual Matches for Fine-Grained Classification. Int. Conf. on Multimedia Retrieval (ICMR), 155-162, 2015.
43. P. Valduriez. Data-intensive HPC: opportunities and challenges. Big Data and Extreme-scale computing (BDEC), Barcelona, Spain, 2015 (invited paper).
44. V. Leveau, A. Joly, O. Buisson, P. Valduriez. Recognizing Thousands of Legal Entities through Instance-based Visual Classification, ACM Multimedia Conf., 1029-1032, 2014.
45. J. Liu, V. Silva, E. Pacitti, P. Valduriez, M. Mattoso. Scientific Workflow Partitioning in Multisite Cloud. Euro-par Bigdatacloud Workshop, 105-116, 2014.
46. E. Castanier, C. Jonquet, S. Melzi, P. Larmande, M. Ruiz, P. Valduriez. Semantic Annotation Workflow using Bio-Ontologies. Workshop on Crop Ontology and Phenotyping Data Interoperability, 2014.
47. J. Dias, E. Ogasawara, D. de Oliveira, F. Porto, P. Valduriez, M. Mattoso. Algebraic Dataflows for Big Data Analysis. IEEE Int. Conf. on Big Data, 150-155, 2013.
48. M. Liroz-Gistau, R. Akbarinia, D. Agrawal, E. Pacitti, P. Valduriez. Data Partitioning for Minimizing Transferred Data in MapReduce. Int. Conf. on Data Management in Grid and P2P Systems (Globe), 1-12, 2013.
49. R. Tang, H. Wu, Z. Bao, S. Bressan, P. Valduriez. The Price is Right. Int. Conf. on Databases and Expert Systems Applications (DEXA), 2013.
50. R. Tang, D. Shao, S. Bressan, P. Valduriez. What you Pay for is What you Get. Int. Conf. on Databases and Expert Systems Applications (DEXA), 2013.
51. R. Coletta, E. Castanier, P. Valduriez, C. Frisch. WebSmatch: a tool for Open Data. Second Workshop on Open Data (WOD), 2013.
52. R. Coletta, E. Castanier, P. Valduriez, C. Frisch, DuyHoa Ngo, Z. Bellahsene. Public Data Integration with WebSmatch. First Workshop on Open Data (WOD), 2012.
53. N. Ayat, H. Afsarmanesh, R. Akbarinia, P. Valduriez. Pay-As-You-Go Data Integration Using Functional Dependencies. Int. Cross-Domain Conference and Workshop on Availability, Reliability, and Security (CD-ARES), 375-389, 2012.
54. N. Ayat, H. Afsarmanesh, R. Akbarinia, P. Valduriez. An Uncertain Data Integration System. ODBASE, OTM Conferences, 825-842, 2012.

55. M. Liroz-Gistau, R. Akbarinia, E. Pacitti, F. Porto, P. Valduriez. Dynamic Workload-Based Partitioning for Large-Scale Databases. Int. Conf. on Databases and Expert Systems Applications (DEXA), 183-190, 2012.
56. A. Chirigati, V. Silva, E. Ogasawara, D. Oliveira, J. Dias, F. Porto, P. Valduriez, M. Mattoso. Evaluating Parameter Sweep Workflows in High Performance Computing. First Int. Workshop on Scalable Workflow Enactment Engines and Technologies (SWEET'12), Held in conjunction with SIGMOD, 2012.
57. L. Barguano, V. Muntes-Mulero, P. Valduriez, D. Dominguez-Sal. ParallelGDB: A Parallel Graph Database based on Cache Specialization. Int. Database Engineering & Applications Symposium (IDEAS), 2011.
58. W. K. Dedzoe, P. Lamarre, R. Akbarinia, P. Valduriez. Efficient Early Top-k Query Processing in Overloaded P2P Systems. Int. Conf. on Databases and Expert Systems Applications (DEXA), 210-224, 2011.
59. P. Valduriez. Principles of distributed Data Management in 2020? Int. Conf. on Databases and Expert Systems Applications (DEXA), Keynote paper, 2011.
60. J. Quiané-Ruiz, P. Lamarre, S. Cazalens, P. Valduriez. Scaling Up Query Allocation in the Presence of Autonomous Participants. Int. Conf. on Database Systems for Advanced Applications (DASFAA), 210-224, 2011.
61. W. K. Dedzoe, P. Lamarre, R. Akbarinia, P. Valduriez. ASAP Top-k Query Processing in Unstructured P2P Systems. IEEE Int. Conf. on Peer-to-Peer Computing (IEEE P2P), 2010.
62. E. Cunha de Almeida, G. Sunye, Y. Le Traon, P. Valduriez. A Framework for Testing Peer-to-Peer Systems. IEEE/ACM Int. Conf. on Automated Software Engineering (ASE), 169-170, 2010.
63. V. Muntes-Mulero, N. Martinez-Bazan, J-L. Larriba-Pey, E. Pacitti, P. Valduriez. Graph Partitioning Strategies for Efficient BFS in Shared-Nothing Parallel Systems. Int. Workshop on Graph Databases, 2010.
64. V. Gulisano, R. Jimenez-Peris, M. Patiño-Martinez, P. Valduriez. StreamCloud: A Large Scale Data Streaming System. IEEE Int. Conf. on Distributed Computing Systems (ICDCS), 126-137, 2010.
65. R. Akbarinia, M. Tlili, E. Pacitti, P. Valduriez, A. Lima. Continuous Timestamping for Efficient Replication Management in DHTs. Int. Conf. on Data Management in Grid and P2P Systems (Globe), LNCS 6265, 38–49, 2010. Selected among the (two) best papers as extended version for the Journal of Transactions on Large-Scale Data and Knowledge-Centered Systems.
66. E. Ogasawara, J. Dias, D. Oliveira, C. Rodrigues, C. Pivotto, R. Antas, V. Braganholo, P. Valduriez, M. Mattoso. A P2P Approach to Many Tasks Computing for Scientific Workflows. Int. Conf. on High Performance Computing for Computational Science (VecPar), 2010.
67. M. Tlili, R. Akbarinia, E. Pacitti, P. Valduriez. Scalable P2P Reconciliation Infrastructure for Collaborative Text Editing. Int. Conf. on Advances in Databases, Knowledge, and Data Applications (DBKDA), IEEE Computer Society, 155-164, 2010.
68. T. Sarni, A. Queudet, P. Valduriez. Real-Time Support for Software Transactional Memory. IEEE Int. Conf. on Embedded and Real-Time Computing Systems and Applications (RTCSA), 477-485, 2009.
69. T. Sarni, A. Queudet, P. Valduriez. Software Transactional Memory: worst case execution time. Int. Conf. on Real-Time and Network Systems (RTNS), 107-114, 2009.
70. M. Jawad, P. Serrano-Alvarado, P. Valduriez, S. Drapeau. A Data Privacy Service for Structured P2P Systems. IEEE Mexican Int. Conf. on Computer Science, 45-56.
71. M. Jawad, P. Serrano-Alvarado, P. Valduriez. Protecting Data Privacy in Structured P2P Networks. Inf. Conf. On Data Management in Grid and Peer-to-Peer Systems (Globe), 85-98, 2009.
72. J. Quiané-Ruiz, P. Lamarre, P. Valduriez. SbQA: A Self-Adaptable Query Allocation Process. IEEE Int. Conf. on Data Engineering (ICDE), Demo paper, 1527-1530, 2009.
73. M. Mattoso, E. Pacitti, P. Valduriez, R. Akbarinia, V. Braganholo, A. A.B. Lima. SARAVÁ: data sharing for online communities in P2P. Colloquium of Computation: Brazil / Inria, Cooperations, Advances and Challenges (Colibri), Best paper award, 2009.

74. W. Palma, R. Akbarinia, E. Pacitti, P. Valduriez. Efficient DHTJoin: Processing Continous Join Queries using DHT Networks. Int. Workshop on Data Management in P2P Systems (DAMAP), co-located with EDBT, 2009.
75. K. Park, P. Valduriez, H. Choo. Mobile Continuous Nearest Neighbor Queries on Air. ACM SIGSPATIAL Int. Conf. on Advances in Geographic Information Systems (GIS), 65, 2008.
76. E. Cunha de Almeida, G. Sunye, Y. Le Traon, P. Valduriez. A Framework for Testing Peer-to-peer Systems. IEEE Int. Conf. on Software Reliability Engineering (ISSRE), 167-176, 2008, Selected as second best paper for Empiral Software Engineering Journal.
77. J. Manjarrez, J. Martinez, P. Valduriez. Efficient Processing of Nearest Neighbor Queries in Parallel Multimedia Databases. Int. Conf. on Databases and Expert Systems Applications (DEXA), 326-339, 2008.
78. E. Cunha de Almeida, G. Sunye, Y. Le Traon, P. Valduriez. Testing Peers' Volatility. IEEE/ACM Int. Conf. on Automated Software Engineering (ASE), 419-422, 2008.
79. M. Jawad, P. Serrano-Alvarado, P. Valduriez. Design of PriServ, a privacy service for DHTs. Int. Workshop on Privacy and Anonymity in Information Society, PAIS), 21-25, 2008.
80. M. Tlili, W. K. Dedzoe, E. Pacitti, P. Valduriez, R. Akbarinia, et al. Estampillage et Journalisation P2P pour Xwiki. Conférence Internationale sur les NOuvelles TEchnologies de la REpartition (NOTERE), ACM, Lyon, 197-200, 2008.
81. W. Palma, R. Akbarinia, E. Pacitti, P. Valduriez. Efficient Processing of Continuous Join Queries using Distributed Hash Tables. European Conf. on Parallel and Distributed Computing (Euro-Par), 632-641, 2008.
82. E. Cunha de Almeida, G. Sunye, P. Valduriez. Testing Architectures for Large Scale Grids. Int. Workshop on High-performance Data Management in GRIDs (HPDgrid), VecPar, LNCS 5336, Springer, 555-566, 2008.
83. D. X. de Sousa, S. Lifschitz, P. Valduriez. BLAST Distributed Execution on Partitioned Databases with Primary Fragments. Int. Workshop on High-performance Data Management in GRIDs (HPDgrid), VecPar, LNCS 5336, Springer, 544-554, 2008.
84. A. Ventresque, S. Cazalens, P. Lamarre, P. Valduriez. Improving Interoperability Using Query Interpretation in Semantic Vector Spaces. 5th European Semantic Web Conference (ESWC), Nominated for best paper (top 4), 539-553, 2008.
85. A. Ventresque, S. Cazalens, P. Lamarre, P. Valduriez. Dealing with P2P Semantic Heterogeneity through Query Expansion and Interpretation. Int. Workshop on Data Management in P2P Systems (DAMAP), co-located with EDBT, 2008.
86. J. Quiané-Ruiz, P. Lamarre, S. Cazalens, P. Valduriez. Managing Virtual Money for Satisfaction and Scale Up in P2P Systems. Int. Workshop on Data Management in P2P Systems (DAMAP), co-located with EDBT, 2008.
87. E. Cunha de Almeida, G. Sunye, P. Valduriez. Action Synchronization in P2P System Testing. Int. Workshop on Data Management in P2P Systems (DAMAP), co-located with EDBT, 2008.
88. R. Hayek, G. Raschia, P. Valduriez, N. Mouaddib. Summary Management in P2P Systems. EDBT Conf., 16-25, 2008.
89. M. Paes, A. Lima, P. Valduriez, M. Mattoso. High-performance Query Processing of a Real-world OLAP Database with ParGRES. Int. Conf. on High Performance Computing for Computational Science (VecPar), Best paper award, LNCS 5336, Springer, 192-205, 2008.
90. R. Akbarinia, E. Pacitti, P. Valduriez. Best Position Algorithms for Top-k Queries. Int. Conf. on Very Large Databases (VLDB), Vienna, Austria, 495-506, 2007.
91. J. Quiané-Ruiz, P. Lamarre, P. Valduriez. SQLB: A Query Allocation Framework for Autonomous Consumers and Providers. Int. Conf. on Very Large Databases (VLDB), Vienna, Austria, 2007, 974-985.
92. A. Ventresque, S. Cazalens, P. Lamarre, P. Valduriez. Query Expansion and Interpretation to Go Beyond Semantic P2P Interoperability. Int. Conf. on Ontologies, DataBases, and Applications of Semantics (ODBASE), 870-877, 2007.

93. R. Akbarinia, E. Pacitti, P. Valduriez. Processing Top-k Queries in Distributed Hash Tables. European Conf. on Parallel and Distributed Computing (Euro-Par), 489-502, 2007.
94. J. Quiané-Ruiz, S. Cazalens, P. Lamarre, P. Valduriez. Satisfaction in Distributed Information Systems: a Mediation Process Case Study. ACM Inf. Conf. on Information and Knowledge Management (CIKM), Lisbon, Portugal, 2007, 947-950.
95. D. Faye, G. Nachouki, P. Valduriez: Semantic Query Routing in SenPeer, a P2P Data Management System. Int. Conf. on Network-Based Information Systems (NBiS), Regensburg, Germany, LNCS 4658 Springer 2007, 365-374.
96. J. Manjarrez Sanchez, J. Martinez, P. Valduriez. On the Usage of Clustering for Content Based Image Retrieval. Int. Computer Science Symposium in Russia (CSR), LNCS 4649, Springer, 281-289, 2007.
97. J. Manjarrez Sanchez, J. Martinez, P. Valduriez. A Data Allocation Method for Efficient Content-Based Retrieval in Parallel Multimedia Databases. Int. Workshop on Parallel and Distributed Multimedia Computing (ParDMCom), LNCS 4743, Springer, 285-294, 2007.
98. R. Hayek, G. Raschia, P. Valduriez, N. Mouaddib. Design of PeerSum: A Summary Service for P2P Applications. Int. Conf. on Advances in Grid and Pervasive Computing (GPC), Paris, LNCS 4459 Springer, 13-26, 2007.
99. R. Akbarinia, E. Pacitti, P. Valduriez. Data Currency in Replicated DHTs. ACM SIGMOD Int. Conf. on Management of Data, Beijing, China, 211-222, 2007.
100. J. Quiané-Ruiz, P. Lamarre, P. Valduriez. KnBest - A Balanced Request Allocation Method for Distributed Information. Int. Conf. on Database Systems for Advanced Applications (DASFAA), Bangkok, Thailand, 237-248, 2007.
101. M. Didonet Del Fabro, Patrick Valduriez. Semi-automatic Model Integration using Matching Transformations and Weaving Models. ACM Symp. On Applied Computing (SAC'07), Model Transformation Track, Seoul, Korea, 963-970, 2007.
102. Y. Kurtev, J. Bézivin, F. Jouault, P. Valduriez. Model-based DSL Frameworks. ACM SIGPLAN Conf. on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA'06), Onward! track, Portland, USA, 602-616, 2006.
103. A. Jouault, F. Allilaire, J. Bézivin, I. Kurtev, P. Valduriez. ATL: a QVT-like transformation language. ACM SIGPLAN Conf. on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA'06), Demo. track, Portland, USA, 719-720, 2006.
104. M. Didonet Del Fabro, J. Bézivin, Patrick Valduriez. Model-driven Tool Interoperability: an Application in Bug Tracking. Int. Conf. on Ontologies, DataBases, and Applications of Semantics (ODBASE'06), Montpellier, 863-881, 2006.
105. J. Quiané-Ruiz, P. Lamarre, P. Valduriez. Satisfaction Based Query Load Balancing. Int. Conf. on Cooperative Information Systems (CoopIS'06), Montpellier, 36-53, 2006.
106. V. Martins, R. Akbarinia, E. Pacitti, P. Valduriez. Reconciliation in the APPA P2P System. IEEE Int. Conf. on Parallel and Distributed Systems (ICPADS), Minneapolis, USA, 401-410, 2006.
107. R. Akbarinia, V. Martins, E. Pacitti, P. Valduriez. Top-k Query Processing in the APPA P2P System. Int. Conf. on High Performance Computing for Computational Science (VecPar 2006), LNCS 4395, Springer, 158-171, 2007.
108. A. Miranda, A. Lima, P. Valduriez, M. Mattoso. Apuama: Combining Intra-query and Inter-query Parallelism in a Database Cluster. Currents Trends in Database Technology – EDBT 2006, LNCS 4254, Springer, 649-661, 2006.
109. V. Martins, E. Pacitti, P. Valduriez. A Dynamic Distributed Algorithm for Semantic Reconciliation. Distributed Data & Structures 6 (WDAS). Records of the 7th Int. Meeting (Santa Clara, California), Waterloo. Carleton Scientific, 2006.
110. A. Furtado, A. Lima, E. Pacitti, P. Valduriez, M. Mattoso. Physical and Virtual Partitioning in OLAP Database Clusters. Int. Symp. on Computer Architecture and High Performance Computing (SBAC), Rio de Janeiro, Brazil, 143-150, 2005.

111. P. Valduriez E. Pacitti, C. Coulon. Large-scale Experimentation with Preventive Replication in a Database Cluster. VLDB Workshop on Design, Deployment and Implementation of Database Replication, Trondheim, Norway, 2005.
112. C. Coulon, E. Pacitti, P. Valduriez. Consistency Management for Partial Replication in a High Performance Database Cluster. IEEE Int. Conf. on Parallel and Distributed Systems (ICPADS), Fukuoka, Japan, 809-815, 2005.
113. J. Bézivin, F. Jouault, P. Rosenthal, P. Valduriez. Modeling in the Large and Modeling in the Small. Model Driven Architecture, European MDA Workshops: Foundations and Applications 2004, LNCS 3599, Springer 33-46, 2004.
114. J. Bézivin, F. Jouault, P. Valduriez. First Experiments with a ModelWeaver. OOPSLA & GPCE Workshop on Best Practices for Model Driven Software Development, Vancouver, Canada, 2004.
115. J. Bézivin, F. Jouault, P. Valduriez. On the Need for Megamodels. OOPSLA & GPCE Workshop on Best Practices for Model Driven Software Development, Vancouver, Canada, 2004.
116. P. Lamarre, S. Cazalens, S. Lemp, P. Valduriez. A Flexible Mediation Process for Large Distributed Information Systems. Int. Conf. on Cooperative Information Systems (CoopIS'04), Cyprus, 19-36, 2004.
117. C. Lepape, S. Gançarski, P. Valduriez. REFRESCO. Improving Query Performance through Freshness Control in a Database Cluster. Int. Conf. on Cooperative Information Systems (CoopIS'04), Cyprus, 174-193, 2004.
118. A. Lima, M. Mattoso, P. Valduriez. Adaptive Virtual Partitioning for OLAP Query Processing in a Database Cluster. Brazilian Symposium on Databases (SBBD), Brazilia, Brazil, 92-105, 2004.
119. A. Lima, M. Mattoso, P. Valduriez. OLAP Query Processing in a Database Cluster. European Conf. on Parallel and Distributed Computing (Euro-Par 2004), Pisa, Italy, 355-362, 2004.
120. P. Valduriez, E. Pacitti. Data Management in Large-scale P2P Systems. Int. Conf. on High Performance Computing for Computational Science (VecPar -2004), Invited paper, Valencia, Spain, 2004. LNCS 3402, Springer, 109-122, 2005.
121. C. Coulon, E. Pacitti, P. Valduriez. Scaling up the Preventive Replication of Autonomous Databases in Cluster Systems. Int. Conf. on High Performance Computing for Computational Science (VecPar -2004), Valencia, Spain, 2004. LNCS 3402, Springer, 174-188, 2005.
122. R. Akbarinia, V. Martins, E. Pacitti, P. Valduriez. Replication and Query Processing in the APPA Data Management System. Distributed Data & Structures 6 (WDAS). Records of the 6th Int. Meeting (Lausanne, Switzerland), Waterloo. Carleton Scientific, 2004.
123. F. Porto, E. Laber, P. Valduriez. Cherry Picking: a semantic query processing strategy for the evaluation of expensive predicates. Brazilian Symposium on Databases (SBBD), Nominated for best award (top 3), Manaus, Brazil, 356-370, 2003.
124. S. Gançarski, C. Lepape, P. Valduriez. Relaxing Freshness for Improving Performance in a Cluster of Replicated Databases. Distributed Data & Structures 5 (WDAS). Records of the 5th Int. Meeting (Thessaloniki, Greece), Waterloo. Carleton Scientific, 2003, 45-60.
125. C. Lepape, S. Gançarski, P. Valduriez. Trading Freshness for Performance in a Cluster of Replicated Databases. Int. Conf. on Cooperative Information Systems (CoopIS'03), Catania, Italy, 14-15, 2003.
126. S. Gançarski, H. Naacke, E. Pacitti, P. Valduriez. Parallel Processing with Autonomous Databases in a Cluster System. Int. Conf. on Cooperative Information Systems (CoopIS'02), Los Angeles, California, 410-428, 2002.
127. H. Naacke, S. Gancarski, P. Valduriez. Load Balancing of Autonomous Databases and Applications in a Cluster system. Distributed Data & Structures 4 (WDAS). Records of the 4th Int. Meeting (Paris). Waterloo. Carleton Scientific, 159-170, 2002.
128. M. Andrei, P. Valduriez. User-Optimizer Communication using Abstract Plans in Sybase ASE. Int. Conf. on VLDB, Roma, 29-38, 2001.
129. N. Anciaux, C. Bobineau, L. Bouganim, P. Pucheral, P. Valduriez. PicoDBMS. Validation and Experience. Int. Conf. on VLDB, Roma, 709-710, 2001.

130. L. Bouganim, F. Fabret, F. Porto, P. Valduriez . Processing Queries with Expensive Functions and Large Objects in Distributed Mediator Systems. Int. Conf. on Data Engineering, Heidelberg, 91-98, 2001.
131. C. Bobineau, L. Bouganim, P. Pucheral, P. Valduriez. PicoDBMS: Scaling down database techniques for the Smartcard. Int. Conf. on VLDB, Best paper award, Cairo, 11-20, 2000.
132. K. Yagoub, D. Florescu, V. Issarny, P. Valduriez. Caching Strategies for Data-intensive Web Sites. Int. Conf. on VLDB, Cairo, 188-199, 2000.
133. K. Yagoub, D. Florescu, V. Issarny, P. Valduriez. Building and Customizing Data-Intensive Web Sites Using Weave. Int. Conf. on VLDB, Cairo, 607-610, 2000.
134. L. Bouganim, F. Fabret, C. Mohan, P. Valduriez. Dynamic Query Scheduling in Data Integration Systems. Int. Conf. on Data Engineering, San Diego, California, 425-434, 2000.
135. L. Bouganim, O. Kapiskaia, P. Valduriez. Memory-Adaptive Scheduling for Large Query Execution. ACM Int. Conf. on Information and Knowledge Management (CIKM), Bethesda, Maryland, 105-115, 1998.
136. P. Valduriez . Parallel Database Systems. Colloque Africain de la Recherche en Informatique (CARI). Dakar, Sénégal, 1998.
137. M. Skubiszewski, P. Valduriez. Concurrent Garbage Collection in O2. Int. Conf. on VLDB, Athens, 356-365, 1997.
138. L. Bouganim, D. Florescu, P. Valduriez. Dynamic Load Balancing in Hierarchical Parallel Database Systems. Int. Conf. on VLDB, Bombay, 436-447, 1996.
139. L. Bouganim, B. Dageville, P. Valduriez. Adaptive Parallel Query Execution in DBS3. Int. Conf. on Extending Database technology, Avignon, 481-484, 1996.
140. D. Florescu, L. Raschid, P. Valduriez. Answering Queries Using OQL View Expressions. Int. Workshop on Materialized Views, in cooperation with ACM SIGMOD, Montreal, 84-90, 1996.
141. A. Tomasic, L. Raschid, P. Valduriez. Scaling Heterogeneous Database and the Design of Disco. IEEE Int. Conf. on Distributed Computing Systems (ICDCS), Hong Kong, 449-457, 1996.
142. D. Florescu, L. Raschid, P. Valduriez. Using Heterogeneous Equivalences for Query Simplification in MDBMS. Int. Conf. on Cooperative Information Systems (CoopIS), Vienna, 158-169, 1995.
143. L. Daynès, O. Gruber, P. Valduriez. Locking in OODBMS Clients Supporting Multiple Nested Transactions. Int. Conf. on Data Engineering (ICDE), Taipei, Taiwan, 316-323, 1995.
144. D. Florescu, P. Valduriez. Rule-based Query Processing in the IDEA System. Int. Symp. on Advanced Database Technologies and Their Integration, Nara, Japan, 1994.
145. M. Novak, G. Gardarin, P. Valduriez. Flora. a Functional-Style Language for Object and Relational Algebra. Int. Conf. on Databases and Expert Systems Applications (DEXA), Athens, 37-46, 1994.
146. O. Gruber, P. Valduriez. Towards Persistent Object Systems for Desktop Computing. Int. Conf. on Cooperative Information Systems (CoopIS), Toronto, 60-71, 1994.
147. J.W. Schmidt, F. Matthes, P. Valduriez. Building Persistent Application Systems in Fully Integrated Data Environments. Modularization, Abstraction and Interoperability. Euro-Arch'93 Congress, Munich, 1993.
148. R. Lanzelotte, P. Valduriez, M. Zait. On the Effectiveness of Optimization Search Strategies for Parallel Execution Spaces. Int. Conf. on VLDB, Dublin, 493-504, 1993.
149. P. Valduriez. Parallel Database Systems. the case for shared-something. Keynote address, IEEE Int. Conf. on Data Engineering, Vienna, Austria, 460-465, 1993.
150. S. Parker, E. Simon, P. Valduriez. SVP, a Data Model Capturing Sets, Streams and Parallelism. Int. Conf. on VLDB, Vancouver, 115-126, 1992.
151. E. Amiel, M-J. Bellosta, P. Valduriez. A Model-Independent Object Schema manager. Int. Workshop on Persistent Object Systems, Pisa, Italy, 224-240, 1992.
152. M-J. Bellosta, P. Valduriez, F. Viallet. View Integration in OODB. Int. Conf. on Database and Expert Systems Applications (DEXA), Valence, Spain, 288-293, 1992.
153. D. Shasha, E. Simon, P. Valduriez. Simple Rational Guidance for Chopping Up Transactions. ACM SIGMOD Int. Conf., San Diego, 298-307, 1992.

154. R. Lanzelotte, P. Valduriez, M. Zait. Optimization of Object-Oriented Recursive Queries Using Cost-Controlled Strategies. ACM SIGMOD Int. Conf., San Diego, 256-265, 1992.
155. G. Gardarin, P. Valduriez. ESQL2, an Extended SQL2 with F-logic Semantics. IEEE Int. Conf. on Data Engineering, Phoenix, Arizona, 320-327, 1992.
156. O. Gruber, L. Amsaleg, L. Daynes, P. Valduriez. Eos, an Environment for Object-based Systems. Hawaii Int. Conf. on System Sciences, Honolulu, 1992.
157. M-J. Bellosta, P. Valduriez, F. Viallet. Reusing Engineering Components in OODB. Int. Conf. on Data and Knowledge Systems for Manufacturing and Engineering, Lyon, 1992.
158. B. Bergsten, M. Couprie, P. Valduriez. Prototyping DBS3, a Shared-Memory Parallel Database System. Int. Conf. on Parallel and Distributed Information Systems (PDIS), Miami, Florida, 226-234, 1991.
159. R. Lanzelotte, P. Valduriez, M. Ziane, J-P. Cheiney. Optimization of Non-Recursive Queries in OODB. Int. Conf. on Deductive and Object-Oriented Databases (DOOD), Munich, 1-21, 1991.
160. R. Lanzelotte, P. Valduriez. Extending the Search Strategy in a Query Optimizer. Int. Conf. on VLDB, Barcelona, 363-373, 1991.
161. M-J. Bellosta, P. Valduriez, F. Viallet. OMNIS, an Object Management Interface System. TOOLS Int. Conf., Paris, 1991.
162. G. Gardarin, P. Valduriez. ESQL, an Extended SQL with Deductive and Object Capabilities. Int. Conf. on Database and Expert System Applications (DEXA), Vienna, 299-306, 1990.
163. P. Pucheral, J-M. Thevenin, P. Valduriez. Efficient Mayn Memory Data Management using the DBGRAPH Storage Model. Int. Conf. on VLDB, Brisbane, Australia, 683-695, 1990.
164. P. Valduriez. Query Processing in the EDS Parallel Database System. Invited paper, Brazilian Symposium on Databases, Rio de Janeiro, Brazil, 1990.
165. P. Valduriez, S. Danforth. Query Optimization in Database Programming Languages. Int. Conf. on Deductive and Object-Oriented Databases, Kyoto, 553-571, 1989.
166. P. Valduriez, S. Danforth, T. Briggs, B. Hart, M. Cochinwala. Compiling FAD, a Database Programming Language. Int. Workshop on Database Programming Languages (DBPL), Portland, Oregon, 375-393, 1989.
167. S. Khoshafian, P. Valduriez. A Parallel Container Model for Data Intensive Applications. Int. Workshop on Database Machines (IWDM), Deauville, 156-170, 1989.
168. B. Hart, S. Danforth, P. Valduriez. Parallelizing FAD, a Database Programming Language. Int. Symp. on Databases in Parallel and Distributed Systems, Austin, Texas, 72-79, 1988.
169. S. Khoshafian, P. Valduriez, G. Copeland. Parallel Query Processing of Complex Objects. IEEE Int. Conf. on Data Engineering (ICDE), Los Angeles, 202-209, 1988.
170. P. Valduriez, S. Khoshafian. Transitive Closure of Transitively Closed Relations. Expert Database Conf., Vienna, Virginia, 377-400, 1988.
171. F. Bancilhon, T. Briggs, S. Khoshafian, P. Valduriez. FAD, a Simple and Powerful Database Language. Int. Conf. on VLDB, Brighton, England, 97-105, 1987.
172. S. Khoshafian, P. Valduriez, G. Copeland. A Query Processing Strategy for the Decomposed Storage Model. IEEE Int. Conf. on Data Engineering (ICDE), Los Angeles, 636-643, 1987.
173. S. Khoshafian, P. Valduriez. Sharing, Persistence, and Object-Orientation: A Database Perspective. Int. Workshop on Database Programming Languages (DBPL), Roscoff, France, 221-240, 1987.
174. S. Khoshafian, P. Valduriez. Parallel Execution Strategies for Declustered Databases. Int. Workshop on Database Machines (IWDM), Tokyo, Japan, 458-471, 1987.
175. P. Valduriez, S. Khoshafian, G. Copeland. Implementation Techniques of Complex Objects. Int. Conf. on VLDB, Kyoto, Japan, 101-110, 1986.
176. G. Copeland, S. Khoshafian, M. Smith, P. Valduriez. Buffering Schemes for Permanent Data. IEEE Int. Conf. on Data Engineering (ICDE), Los Angeles, 214-221, 1986.
177. P. Valduriez, H. Boral. Evaluation of Recursive Queries Using Join Indices. Expert Database Conf., Charleston, South Carolina, 271-293, 1986.

178. E. Simon, P. Valduriez. Integrity Control in Distributed Database Systems. Hawaii Int. Conf. on System Sciences, Honolulu, 1986.
179. P. Faudemay, P. Valduriez. Design and Analysis of a Direct Filter Using Parallel Comparators. Int. Workshop on Database Machines (IWDM), Bahamas, 137-152, 1985.
180. E. Simon, P. Valduriez. Design and Implementation of an Integrity Subsystem. ACM SIGMOD Int. Conf., Boston, 9-17, 1984.
181. P. Valduriez, Y. Viemont. A New Hashing Scheme using Predicate Trees. ACM SIGMOD Int. Conf., Boston, 107-114, 1984.
182. P. Valduriez, L. Verlaine. Design of a Database Machine Using Conventional and Specialized Hardware. IEEE Trends and Application Conf., Gaithersburg, MD, 1984.
183. M. Bouzeghoub, P. Valduriez. New Tools for Logical Database Design. BCS Int. Conf. on Databases, London, 1984.
184. G. Gardarin, P. Valduriez, Y. Viemont. Predicate Trees. A Way of Optimizing Relational Queries. IEEE Int. Conf. on Data Engineering (ICDE), Los Angeles, 439-444, 1984.
185. P. Valduriez. Design of a Tightly Coupled Database Machine. Int. Seminar on Distributed Data Sharing, Parma, Italy, 1984.
186. G. Gardarin, P. Bernadat, P. Valduriez, Y. Viemont. Design of a Multiprocessor Relational Database System. IFIP World Computer Congress, Paris, 1983.
187. P. Valduriez. Semi-join Algorithms for Multiprocessor Systems. ACM SIGMOD Int. Conf., Orlando, Florida, 225-233, 1982.
188. P. Valduriez. Semi-join Algorithms for Distributed Database Machines. Int. Symp. on Distributed Data Bases, Berlin, 1982.
189. G. Gardarin, P. Valduriez. The SABRE Database Machine. Int. Conf. on Management of Distributed Data Processing, Paris, 1982.

+ over 100 papers published in national journals and conferences, in particular the French BDA conference.

Misc.

- P. Valduriez. Making the Right Move to Senior Researcher. ACM SIGMOD Record, May 2021.
- G. Weikum, P. Valduriez. Report on the 2004 SIGMOD Conference. ACM SIGMOD Record, 33(4), 2004.
- P. Valduriez. Hints to Write Technical Papers. Engineering of Information Systems, Hermes, 2(3), 1994.

Theses

- P. Valduriez. Optimization of relational operators in database machines. Doctorat d'Etat in computer science, UPMC, Paris, 20 Sept. 1985.
- P. Valduriez. Design and implementation of a join and sort processor for a database machine.Ph.D. thesis in computer science, UPMC, Paris, 16 Dec. 1981.