

Alexander Russakovsky, Ph.D.

3828 Corina Way
Palo Alto, CA 94303
Email:alex@russakovsky.com

Tel. (650) 843-0737
Fax (650) 384-0305
Mobile (650) 464-1306

Skills:

- Exceptional combination of very strong mathematical and software development background
- Focus on solving hard problems and finding the right approaches.
- 25+ years of software development experience
- Database internals: data storage models, indexing, query execution, calculation engine
- Distributed databases and applications; multiple inventions and patents
- Big Data analytics, business intelligence, decision support, data integration and virtualization
- OLAP, forecasting, data mining, streaming, modeling, metadata, universal data access, semantic layer
- Data modeling: relational, multidimensional and heterogeneous schemas, indexing, etc.
- Software architecture, object oriented analysis, design and programming.
- Java programming on all platforms (from the moment Java appeared), J2SE, J2EE, etc.
- C, C++, XML, etc.
- Complex algorithms and data structures
- Data analysis: statistics, stochastic processes, data mining, machine learning, risk management.
- Numerical methods, mathematical programming, financial models and algorithms.
- Articles in computer journals, finance, white papers, research articles.
- Member of XML for Analysis, Java OLAP (JCP JSR-69) and Java Data Mining (JCP JSR-73, JCP JSR-247) expert groups working on industry standards
- Speaker at Java One conference 2001, 2002; seminar talks at Bell Labs, IBM Silicon Valley Labs, Stanford University DB Group; Telecom Council Big Data Analytics panelist 2014, 2016
- Participant of VLDB conferences 2001, 2004, 2005, 2006, 2012, 2013, ICDT 2008, EDBT 2014

Selected positions in Software Development:

2023- **Zeepabyte, CTO and Founder**

- Invention and implementation of Vector Database technology (Approximate Nearest Neighbor Search) based on Zeepabyte Cascade (Java + Python integration)

2021-2023 **Oracle Corporation, Architect – Oracle Analytics Cloud**

- Various performance and usability improvement proposals for Oracle BI Server and OAC
- Performance improvement for Parquet file export in OBIS (C++)
- Performance improvement for Parquet file export in Dataset Services (Java)
- Project for offloading OAC user data sets to external DB engines (ADW and Heatwave)
- Participation in Oracle patent committee

2020-2021 **Workday, Architect – Planning and Multidimensional (Consultant)**

- Assessment of Unified Reporting project between Planning and Analytics teams
- Analysis, assessment and recommendations for improving scalability and performance of Planning Evaluator Engine, numerous presentations and proposals

2015-2020

Zeepabyte, CTO and Founder

- Invention and implementation of unique multidimensional indexing technology, Cascade
- Design and development of Cascade Analytics System for Big Data BI and IoT (Java)
- Record breaking hyperscale performance in Star Schema Benchmark with Mellanox Technologies Labs: 2x faster than industry leaders on 10x less hardware.
- Patent application granted
- SQL, Natural Language and Predictive query support
- Built in distributed Machine Learning framework and algorithms
- Automatic ETL execution plan generation from declared input schema
- Support for S-3 connectors
- Numerous customer POCs

2011-2015

Futurewei Technologies (Huawei USA R&D Labs), Technical VP

- Chief architect and contributor of Business Intelligence Platform for Telecom project (Java)
- Invention and implementation of new unique Grasshopper algorithms for fast scanning of multidimensional data (patent application filed and paper written - see <http://arxiv.org/abs/1310.0141>)
- Design and development of Distributed Multidimensional Database as the foundation of Data Services within the BI platform. Works with multiple back-ends including in-memory and disk based, distributed and standalone. Performs on Big Data orders of magnitude faster than all tested competitors
- High performance ad hoc query benchmark winner with Hadoop/HBase backend (Big Data MOLAP)

2009-2011

Oracle Corporation, Architect

- OBIEE Platform architecture with focus on multidimensional, metadata, streaming and ETL aspects
- BI Platform automation, data introspection, Predictive BI and data mining integration
- Development of new features for the Essbase Calculator Framework
- Numerous project proposals and presentations
- Evaluation of development initiatives
- Participation in Oracle patent committee

2007-2009

Truviso Inc, VP and Chief Scientist, Analytics

- Theory, interface and design for streaming DB engine language and implementation (Postgres, C). In particular, stream-stream joins, non-contiguous time concept design and implementation, etc.
- Workbench for evaluating financial algorithms and strategies through treating them as classification algorithms and comparing performance.
- BI/analytics/data mining strategy and implementation.
- Metadata modeling, query generation and application building tools.
- Streaming OLAP, theory and conceptual design, materialized view selection, etc.
- Numerous customer POCs.

2007

Oracle Corporation, Architect

- Oracle acquired Hyperion...

1998 - 2007

Hyperion Solutions Corporation, Distinguished Software Engineer / Architect

- Design of the metadata layer, dependency and lineage analysis, efficient query generation and execution in a heterogeneous database environment (SQL, OLAP and other sources) based on logical model within the enterprise analytical platform

- Design of the next generation enterprise analytical platform (new data storage model and kernel, new multidimensional indexing, new data and metadata model, new language, new calculation model, task management, new distributed architecture, etc) (managed a team of about 15 people). In particular, the new kernel released in Essbase.7x improved performance and scalability hundreds of times and is now the world's leading multidimensional database engine by a significant margin.
- Design and implementation of the MDX/XML for Analysis engine within the Essbase server
- Design and implementation of Java-based forecasting and data mining framework for Essbase (patent awarded). The framework pioneers data mining within a multidimensional database in a commercial product. Includes regression, clustering, neural networks, decision tree, association rules and naïve Bayes algorithms. See <http://cs.russakovsky.com/dmmd.pdf>
- Design and development of calculator framework for Essbase multidimensional database server with extensible plug-in architecture, including a preprocessor, a parser, an optimizer and an executor (C)
- Research and development of forecasting, approximation and statistical routines for the Essbase calculator (C)
- Design of the joint XML for Analysis specification of Hyperion and Microsoft; headed Hyperion technical team
- Design of globalization (Unicode based) architecture for the Essbase platform
- Design and development of user defined functions and macros framework for Essbase using embedded Java Virtual Machine
- Design and development of Java Virtual Machine attachment to the Essbase server
- Dependency analysis for parallel computations in the Essbase calculator engine
- Research paper on theoretical foundations of multidimensional database models (<http://cs.russakovsky.com/multidimensionaldatamodel.pdf>)
- Research paper on what-if OLAP queries with changing dimensions (<http://cs.russakovsky.com/persVLDB.pdf>, see also <http://www.youtube.com/watch?v=CGBhSmi4euo>)
- Industrial track paper on streaming database technology (<http://cs.russakovsky.com/cidr2009.pdf>)
- Research and design of various aspects of streaming technology, in particular streaming OLAP, BI and analytics
- Research, design and prototype of a unique multidimensional indexing scheme
- Research, design and prototype of compression routines for sparse arrays
- Development of stock analysis routines in Java for Hyperion Investor
- Various documents and white papers on product design, architecture and strategic platform description
- Evaluation and analysis of various projects, technologies and acquisition candidates
- Numerous POCs for streaming analytics and BI

1999 - 2003

Econophysica, Technology Director

- Organization of development of financial instruments pricing and modeling software (C and Java) (www.econophysica.com)
- Research of exotic financial instruments' pricing algorithms (e.g. <http://fin.russakovsky.com/fallbacks.pdf>)

1996 - 1998

C.ATS Software Inc., Member of Technical Staff

- Financial models: research, development, design and implementation
- Development of a C++ and Java class libraries for pricing of foreign exchange options and interest rate derivatives
- FX options: European, American, Barriers, Digitals of all kinds, Lookbacks, Average rate/strike, Compound, Double Barriers
- Methods: PDE models, Black-Scholes, binomial trees, Monte Carlo, etc
- Interest rate models: Hull-White model (with calibration), Black's model, etc
- Index amortizing swaps, Asian caps/floors
- High-precision routines for computing normal and binormal cumulative distributions

- Software architecture, model API, remote models, CORBA, etc

1996 - ... **Numerical Objects, Founder**

- Development of high-performance distributed Java numerical math tools (www.numobjects.com)

1989 - 1995 **Science & Engineering Software Experts, Ltd., Project Lead**

- Supervised a team for design and development of a scientific calculator on OS/2 and Windows platforms. In both versions, the mathematical engine, written in C, was placed in dynamic link libraries while the user interface for Windows was written in C++ and for OS/2 - in Visual REXX.
- Worked in a team to develop PRONTO (Professional Numerical Tools) library of mathematical routines. Designed, developed (in C language) and tested the library containing over 70 routines and several original algorithms, and demonstration programs.
- Headed a team to design and develop a control system for conducting experiments involving signal processing, analysis and management for a research lab at the Institute of Technology. Written in C, the system imported data via modem, performed spectral analysis of the signals using Fast Hartley Transform, constructed the mathematical model of the experiment and provided instructions to the operator.

Education:

- MS in Mathematics, Kharkov National University (Ukraine)
- PhD in Mathematics (Multidimensional Complex Analysis), Kharkov National University

Academic Affiliation History:

Various permanent and visiting professor positions at

- Department of Mathematics, Stanford University, Stanford, USA
- Institut J.Fourier, Université Grenoble, France
- Mathematical Sciences Research Institute, Berkeley, USA
- Department of Mathematics, Johns Hopkins University, Baltimore, USA
- Mathematical Division, Institute for Low Temperature Physics, Kharkov, Ukraine
- Fellow of Alexander von Humboldt Foundation, Germany
- Math. Institute, Heinrich Heine Universität Düsseldorf, Germany
- UFR Mathématiques, Université Paul Sabatier, Toulouse, France
- Mathematical Department, Kharkov National University, Ukraine

CONFERENCES, TALKS, PUBLICATIONS

26 international conferences, 28 scientific publications, 65 invited seminar talks in 13 countries, in particular in the USA: Stanford University (database and math seminars), MSRI/UC Berkeley, Bell Labs - AT&T, Silicon Valley Research Center - IBM, Iterated Systems Inc, Johns Hopkins University - Baltimore, University of Maryland - College Park, University of Wisconsin - Madison, Indiana University - Bloomington, University of Michigan - Ann Arbor, Purdue University, University of Illinois - Urbana, University of Illinois - Chicago, University of Chicago, Cornell University, University of Notre Dame, Ohio State University - Columbus, State University of New York - Albany, State University of New York - Stony Brook, University of Washington - Seattle, University of Southern California - Los Angeles

PUBLICATIONS: <https://scholar.google.com/citations?hl=en&user=GBUleq8AAAAJ>