



# ECRC Communication Meeting

12 October 2021



جامعة الملك عبدالله  
للعلوم والتقنية  
King Abdullah University of  
Science and Technology

Extreme Computing  
Research Center

# Happy Birthday, Lewis



Lewis Frye Richardson

- 11 October 1881 born Newcastle Upon Tyne
- 1903 Natural Sciences Tripos, Cambridge
- 1903-09 National Physical Laboratory
- 1909-12 Sunbeam Lamp Company
- 1912-13 Manchester College
- 1913-16 Meteorological Office
- 1916-19 WW I, in France, ambulance crew
- 1920-29 Westminster College
- 1929-40 Paisley Technology School
- 1940 Retired “so that he could *work*” 😊  
on pacifist manifestos
- 1953 Died in Scotland
  
- 1922 Numerical Weather Prediction book
- 1926 Elected to Royal Society
  
- Gave us: Richardson extrapolation  
Richardson iteration  
fundamental work on fractals  
first practical application of finite  
difference methods in PDEs

# Richardson quotation

Big whorls have little whorls,  
That feed on their velocity;  
And little whorls have lesser whorls,  
And so on to viscosity.

[verse to summarize his 1920 paper  
*The supply of energy from and to Atmospheric Eddies*]

# Richardson quotation

An advantage of a mathematical statement is that it is so definite that it might be definitely wrong; and if it is found to be wrong, there is a plenteous choice of amendments ready in the mathematicians' stock of formulae. Some verbal statements have not this merit; they are so vague that they could hardly be wrong, and are correspondingly useless.

*[Mathematics of War and Foreign Politics]*

# An early visionary: L. F. Richardson



from 1922 book on  
*Weather Prediction by Numerical Processes*

# Richardson's parallel computer

In his 1922 book *Weather Prediction by Numerical Processes*, British fluid dynamicist L. F. Richardson proposed the creation of a global weather prediction facility, which he dubbed the “forecast factory.” It would employ some 64,000 human computers sitting in tiers around the circumference of a gigantic globe to dynamically calculate a PDE model of the global weather.

*“Imagine a large hall like a theatre, except that the circles and galleries go right round through the space usually occupied by the stage. The walls of this chamber are painted to form a map of the globe. The ceiling represents the north polar regions, England is in the gallery, the tropics in the upper circle, Australia on the dress circle and the Antarctic in the pit.*

*A myriad computers are at work upon the weather of the part of the map where each sits, but each computer attends only to one equation or part of an equation. The work of each region is coordinated by an official of higher rank. Numerous little "night signs" display the instantaneous values so that neighbouring computers can read them. Each number is thus displayed in three adjacent zones so as to maintain communication to the North and South on the map.”*

# Richardson's parallel computer

*“From the floor of the pit a tall pillar rises to half the height of the hall. It carries a large pulpit on its top. In this sits the man in charge of the whole theatre; he is surrounded by several assistants and messengers. One of his duties is to maintain a uniform speed of progress in all parts of the globe. In this respect he is like the conductor of an orchestra in which the instruments are slide-rules and calculating machines.*

*Four senior clerks in the central pulpit are collecting the future weather as fast as it is being computed, and despatching it by pneumatic carrier to a quiet room. There it will be coded and telephoned to the radio transmitting station. Messengers carry piles of used computing forms down to a storehouse in the cellar.*

*In a neighbouring building there is a research department, where they invent improvements. But there is much experimenting on a small scale before any change is made in the complex routine of the computing theatre. In a basement an enthusiast is observing eddies in the liquid lining of a huge spinning bowl, but so far the arithmetic proves the better way. In another building are all the usual financial, correspondence and administrative offices. Outside are playing fields, houses, mountains and lakes, for it was thought that those who compute the weather should breathe of it freely.”*

# Happy Birthday, Michael



Michael Stonebraker, 1943-

- Born **11 October** 1943 (Milton, NH, USA)
- 1965, B.S.E., EE, Princeton
- 1971, Ph.D., CS&E, Michigan
  - Thesis: *The Reduction of Large Scale Markov Models for Random Chains*
- 1971-2000, Berkeley
- 2001-, MIT
- Foundational in relational databases
  - Developer: Ingres and Postgres
- Started many companies:
  - Ingres, Illustra, Cohera, StreamBase, Vertica, VoltDB, Tamr, Paradigm4
- Now working on:
  - Aurora, C-Store, H-Store, Morpheus, Goby, SciDB
- CTO: Informix
- Now doing data management systems
- Recognition
  - 1969, IEEE John von Neumann Medal
  - 1994, ACM SIGMOD Innovation Award
  - 1994, Fellow, ACM
  - 1997, National Academy of Engineering
  - 2014, Turing Award

# Ingres (INteractive Graphics REtrieval System)

## Ingres (database)

From Wikipedia, the free encyclopedia



This article **needs additional citations for verification**. Please help [improve this article](#) by [adding citations to reliable sources](#). Unsourced material may be challenged and removed. *(June 2011)* ([Learn how and when to remove this template message](#))

**Ingres Database** (/ɪŋˈɡrɛs/ *ing-GRESS*) is a commercially supported, open-source SQL relational database management system intended to support large commercial and government applications. Ingres Database is fully open source with a global community of contributors. However, **Actian** Corporation controls the development of Ingres and makes certified **binaries** available for download, as well as providing worldwide support.

Ingres began as a research project at the **University of California, Berkeley**, starting in the early 1970s and ending in 1985.<sup>[2]</sup> The original code, like that from other projects at Berkeley, was available at minimal cost under a version of the **BSD license**. Ingres spawned a number of commercial database applications, including **Sybase**, **Microsoft SQL Server**, **NonStop SQL** and a number of others. **Postgres (Post Ingres)**, a project which started in the mid-1980s, later evolved into **PostgreSQL**. It is **ACID** compatible and is fully transactional (including all **DDL** statements) and is part of the **Lisog** open-source stack initiative.

### Ingres

Ingres Corporation logo from 2007

<b>Original author(s)</b>	University of California, Berkeley
<b>Developer(s)</b>	Actian Corporation
<b>Stable release</b>	10.2 / September 30, 2014 <sup>[1]</sup>
<b>Written in</b>	C
<b>Operating system</b>	Cross-platform
<b>Type</b>	RDBMS
<b>License</b>	GNU General Public License or proprietary
<b>Website</b>	<a href="http://www.actian.com/products/operational-databases/ingres/">www.actian.com/products/operational-databases/ingres/</a>

# Berkeley webpage

---

## Michael Stonebraker



---

## Biographical Summary:

Dr. Stonebraker is Professor of Electrical Engineering and Computer Sciences at the University of California at Berkeley, where he joined the faculty in 1971. He is widely recognized as one of the world's foremost experts in database technology and is noted for his insight in operating systems and expert systems.

He received a Bachelor of Science degree from Princeton University and Master of Science and Doctor of Philosophy degrees from the University of Michigan. Dr. Stonebraker has held visiting professorships at the Pontifico Universidade Catholique (PUC), Rio de Janeiro, Brazil; the University of California, Santa Cruz; and the University of Grenoble, France.

Dr. Stonebraker founded Ingres Corp. in 1980 and served on the company's board of directors until 1993. (Ingres Corp. was purchased by Computer Associates in 1994.) INGRES, the company's primary product, was a commercialization of Dr. Stonebraker's [INGRES](#) research project into relational database management systems (RDBMS) at Berkeley. Ingres Corp. has been widely recognized as a leader in RDBMS technology.

More recently, Dr. Stonebraker has managed another research project known as [POSTGRES](#), one of the most advanced such projects ever undertaken on the Berkeley campus. In August 1992, Dr. Stonebraker co-founded [Illustra Information Technologies, Inc.](#) in order to commercialize his research in object-relational database technology. He is a member of the company's board of directors and serves as its Chief Technology Officer.

Dr. Stonebraker is currently focusing on DBMS support for [visualization environments](#) and on [next-generation distributed DBMSs](#). In addition, he leads a project which is developing [alternative data management strategies](#) for NASA's Earth Observing System (EOS).

His professional activities include:

1. General Chairman of SIGMOD annual conference, 1987
2. Member of SIGMOD Advisory Committee, 1989 to present
3. Program Chairman of SIGMOD annual conference, 1992
4. Member of Technical Advisory Committee for Citicorp, DB Software, and Bull

# MIT webpage



MIT COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY

SEARCH

[ABOUT CSAIL](#) [RESEARCH](#) [NEWS + EVENTS](#) [RESOURCES](#) [PEOPLE](#) [ALUMNI & FRIENDS](#)



## PEOPLE

[Principal Investigators](#)

[All Members](#)

[Student Spotlights](#)

[Home](#) » [People](#) » [Michael Stonebraker](#)

## MICHAEL STONEBRAKER



[ [Login to edit profile](#) ]

Position: Adjunct Professor  
Office: [32-G922](#)  
Phone: +1 (617) 253-3538  
Email: [stonebraker@csail.mit.edu](mailto:stonebraker@csail.mit.edu)  
Areas of Study: Systems & Networking  
Last Update: July 9, 2008  
[Download vCard](#)

### BIOGRAPHY

Michael Stonebraker has been a pioneer of data base research and technology for more than a quarter of a century. He was the main architect of the INGRES relational DBMS, the object-relational DBMS, POSTGRES, and the federated data system, Mariposa. All three prototypes were developed at the University of California at Berkeley where Stonebraker was a Professor of Computer Science for twenty five years. He is the founder of three successful Silicon Valley startups, whose objective was to commercialize these prototypes.

Professor Stonebraker is the author of scores of research papers on data base technology, operating systems and the architecture of system software services. He was awarded the prestigious ACM System Software Award in 1992, for his work on INGRES. Additionally, he was awarded the first annual Innovation award by the ACM SIGMOD special interest group in 1994, and has been recognized by Computer Reseller News as one of the top five software developers of the century. Moreover, Forbes magazine named him one of the 8 innovators driving the Silicon Valley wealth explosion during their 80th anniversary edition in 1998. He was elected to the National Academy of Engineering in 1998 and is presently an Adjunct Professor of Computer Science at M.I.T.

### PUBLICATIONS

Stan Zdonik, Michael Stonebraker, Mitch Cherniack, Ugur Cetintemel, "Aurora\*: A Distributed Stream Processing System", to be submitted to CIDR 2003.

Don Carney, Ugur Cetintemel, Mitch Cherniack, Christian Convey, Sangdon Lee, Michael Stonebraker, Nesime Tatbul, Stan Zdonik, "Monitoring Streams: A New Class of Data Management Applications", VLDB 2002 (expanded version will appear in the VLDB Journal).

Michael Stonebraker, "Too Much Middleware", SIGMOD Record, March 2001.

Chris Olston, Allison Woodruff, Alexander Aiken, Michael Chu, Vuk Ercegovic, Mark Lin, Mybrid Spalding, Michael Stonebraker, "DataSplash", SIGMOD 1998, p550-552.

Allison Woodruff, Michael Stonebraker, "Supporting Fine-grained Data Lineage in a Database Visualization Environment", ICDE 1997, p91-102.

Alexander Aiken, Jolly Chen, Michael Stonebraker, Allison Woodruff, "Tioga-2: A Direct Manipulation Database Visualization Environment", ICDE 1996, p208-217.

Michael Stonebraker, Paul M. Aoki, Witold Litwin, Avi Pfeffer, Adam Sah, Jeff Sidell, Carl Staelin, Andrew Yu, "Mariposa: A Wide-Area Distributed Database System", VLDB Journal 5(1), p48-63 (1996).

Michael Stonebraker, Paul M. Aoki, Robert

# Turing announcement

## Database pioneer Stonebraker rocks \$1M "Nobel Prize in Computing"



Credit: [Wikipedia](#)

### A.M. Turing Award recognizes Ingres RDBMS creator



By **Bob Brown** | Follow  
Network World | Mar 25, 2015 7:29 AM PT

#### RELATED TOPICS

- Applications
- Big Data/Business

Michael Stonebraker, whose database software breakthroughs helped to tame information overload long before we referred to it as big data, is the recipient of the [2014 ACM A.M. Turing Award](#),

#### RELATED



Stanford crypto expert Dan Boneh wins \$175K computer science award



Crypto dream team Diffie & Hellman wins \$1M "Nobel Prize of Computing"

2011 timeline of major high tech awards: Cryptographers, Unix pioneers lead...

on [IDG Answers](#) →  
How can Convert OST file data into PST file formats?

# Turing citation



A.M. TURING CENTENARY CELEBRATION WEBCAST

acm  
MORE ACM AWARDS

A.M. TURING AWARD

Search TYPE HERE

A.M. TURING AWARD WINNERS BY...

ALPHABETICAL LISTING YEAR OF THE AWARD RESEARCH SUBJECT



## MICHAEL STONEBRAKER

United States – 2014

### CITATION

For fundamental contributions to the concepts and practices underlying modern database systems.



SHORT ANNOTATED  
BIBLIOGRAPHY



ACM DL  
AUTHOR PROFILE



ACM TURING AWARD  
LECTURE VIDEO



RESEARCH  
SUBJECTS



ADDITIONAL  
MATERIALS

### BIRTH:

October 11, 1943 in Newburyport, Mass.

### EDUCATION:

Bachelor's degree in Electrical Engineering (Princeton University, 1965); M.Sc. in Electrical Engineering (University of Michigan, Ann Arbor, 1967); Ph.D. in Computer Science & Engineering (University of Michigan, Ann Arbor, 1971).

### EXPERIENCE

Michael Stonebraker's contributions to the refinement and spread of database management technology are hard to overstate. He began work in this area as a young assistant professor at the University of California—Berkeley. After reading Edgar F. Codd's seminal papers on the relational model, Stonebraker started work with a colleague, Eugene Wong, to develop an efficient and practical implementation. The result was INGRES, a name that reflected the project's original intention to produce a geographically-oriented system with graphical capabilities. This officially stood for "Interactive Graphic and Retrieval System" but echoed the name of a [celebrated French painter](#).

A prototype of INGRES was working by 1974, but the project did not stop there. Over the next decade INGRES, and systems inspired by it, built a new commercial market of relational database systems. Today the relational database management system is one of computing's most important and widely used technologies, having replaced filing cabinets as the standard way of storing and retrieving information.

# Stonebraker quotation

“Knowing what I know now, I would never have started building INGRES, because it’s too hard.... So I think my advice to my younger self would be to suspend your disbelief and just do it anyway. The way you climb Mt. Everest is one step at a time...”

# Stonebraker quotation

“Go spend some time in the real world and work on problems people want solved... [too many people] are more interested in working on problems that are solvable rather than problems that are important.”

# Plan of today's meeting

- People
  - new members, visitors, interns, departing ?
- Announcements
- Tech Transfer Office presentation
- ECRC Culture conversation
- Roundtable



**Other introductions?**

**By self?**

**By sponsor?**

# TTO Presentation

**Dr. Salahud Din**

**PhD, 2012 Imperial College London, Materials Science & Nanotechnology  
At KAUST since 2017**

- Business Development, University-Industry Collaborations, Research Collaborations, Faculty Consultancy, Technology Piloting, Creation of Startup Companies, Expansion of the KAUST Research & Technology Park (KRTP) and facilitation of IP licensing deals.
- Identifying suitable partners (industry and others) to drive the deployment of innovative products from the intellectual assets of KAUST researchers. Externally focused role and responsible for deploying KAUST technologies to partners culminating in options, licenses and/or assignments of KAUST IP rights. Working closely with other members of the Technology Transfer Office (TTO) to achieve the objectives of the Office of Innovation and Economic Development for the benefit of KAUST and the Kingdom of Saudi Arabia.

# ECRC-originated and curated software

- HiCMA [[github.com/ecrc/hicma](https://github.com/ecrc/hicma)] *Hierarchical Computations on Manycore Architectures*, a dense linear algebra library exploiting data sparsity of the matrix operator
- ExaGeoStat [[github.com/ecrc/exageostat](https://github.com/ecrc/exageostat)] *Exascale Geospatial Statistics*, a parallel high performance framework for large-scale computational geostatistics on many-core and distributed systems
- ExaGeoStatR [[github.com/ecrc/exageostatR](https://github.com/ecrc/exageostatR)] an R-wrapped version of ExaGeoStat, to offer the interactive features of the consortium-driven statistical toolbox R (see [r-project.org](https://r-project.org))
- H2Opus [[github.com/ecrc/h2opus](https://github.com/ecrc/h2opus)] linear algebra based on H2 hierarchical matrices, for CPUs, GPUs, and distributed systems
- KBLAS-CPU [[github.com/ecrc/kblas-cpu](https://github.com/ecrc/kblas-cpu)] *KAUST's Basic Linear Algebra Subroutines*, a high performance library implementing a subset of the standard BLAS as well as LAPACK routines on CPUs
- KBLAS-GPU [[github.com/ecrc/kblas-gpu](https://github.com/ecrc/kblas-gpu)] *KAUST's Basic Linear Algebra Subroutines*, a CUDA version of KBLAS-CPU for NVIDIA GPUs

# ECRC-originated and curated software

- KSVD [[github.com/ecrc/ksvd](https://github.com/ecrc/ksvd)] *KAUST's Singular Value Decomposition* solver, a high performance software framework for computing a dense SVD on manycore shared and distributed memory systems
- AI4SAN [[github.com/ecrc/al4san](https://github.com/ecrc/al4san)] *Abstraction Layer library For Standardizing APIs of task-based eNgin*, a lightweight software library that provides APIs to unify the expression of tasks and their data dependencies from (so far) four of the major dynamic runtime engines: QUARK, PaRSEC, StarPU, and OpenMP-LLVM
- STARS-H [[github.com/ecrc/al4san](https://github.com/ecrc/al4san)] *Software for Testing Accuracy, Reliability and Scalability of Hierarchical* computations, a generation and analysis library for H-matrices
- HCORE [[github.com/ecrc/al4san](https://github.com/ecrc/al4san)] BLAS operations for matrices maintained in tile low-rank format
- GIRIH [[github.com/girih](https://github.com/girih)] multicore wavefront diamond-tiling methods with an autotuning harness for fast stencil computations in manycore shared and distributed memory systems

# ECRC-originated and curated software

- `mpi4py-fft` [[pypi.org/project/mpi4py-fft/](https://pypi.org/project/mpi4py-fft/)] a Python package for computing Fast Fourier Transforms (FFTs) exploiting transpose-in-the-network functionality of MPI
- GRACE [[github.com/sands-lab/grace](https://github.com/sands-lab/grace)] GRAdient ComprEssion, a framework for implementing and evaluating gradient compression techniques for distributed deep learning
- SSDC [[github.com/ecrc/ssdc](https://github.com/ecrc/ssdc)] *Entropy Stable Discontinuous Collocation*, an *hp*-adaptive high performance CFD toolkit based on a general spectral element solver for unstructured grids
- MLBS [[github.com/ecrc/mlbs](https://github.com/ecrc/mlbs)] *Multi-Layer Buffer System*, a generalization of two-level cache protocols to N-level memory, which has been adopted for Aramco seismic image processing
- ExaWave [a deliverable to Aramco] is a home-grown Reverse Time Migration (RTM) seismic inversion code that is used to prototype performance improvements for Aramco's GeoDrive

# ECRC-originated and curated software

- *DAIET Data Aggregation In nETwork*, leveraging the programmable data plane to reduce data traffic as it is being forwarded towards the destination by offloading the aggregation task to the network  
[[github.com/AmedeoSapio/DAIET](https://github.com/AmedeoSapio/DAIET)]
- SwitchML: Scaling Distributed Machine Learning with In-Network Aggregation [github.com/p4lang/p4app-switchML]
- OmniReduce: Efficient Sparse Collective Communication and its application to Accelerate Distributed Deep Learning  
[[github.com/sands-lab/omnireduce](https://github.com/sands-lab/omnireduce)]

# ECRC-curated software

- PETSc [[www.mcs.anl.gov/petsc/](http://www.mcs.anl.gov/petsc/)] *Portable, Extensible Toolkit for Scientific Computing*, an R&D100 Award winning toolkit of distributed data structures and solvers, supported out of Argonne National Laboratory, a workhorse code of the US DOE
- UG4 [[github.com/ug4](https://github.com/ug4)] a simulation framework for modeling differential equations on unstructured finite element grids
- SPECFEM3D [[github.com/geodynamics/specfem3d](https://github.com/geodynamics/specfem3d)] *Spectral Finite Element Method in 3D*, a modeling tool for acoustic, elastic, coupled acoustic/elastic, poroelastic, or seismic wave propagation in any type of conforming mesh of hexahedra, structure or unstructured
- CLAWPACK [[clawpack.org](http://clawpack.org)] *Conservation Laws Package*, a toolkit for linear and nonlinear hyperbolic systems of conservation laws, implementing high-resolution Godunov type methods using limiters
- PyCLAW [[clawpack.org/pyclaw/](http://clawpack.org/pyclaw/)] a Python-based interface to the algorithms of Clawpack and SharpClaw, and

# ECRC-curated software

- PETSc [[www.mcs.anl.gov/petsc/](http://www.mcs.anl.gov/petsc/)] *Portable, Extensible Toolkit for Scientific Computing*, an R&D100 Award winning toolkit of distributed data structures and solvers, supported out of Argonne National Laboratory, a workhorse code of the US DOE
- UG4 [[github.com/ug4](https://github.com/ug4)] a simulation framework for modeling differential equations on unstructured finite element grids
- SPECFEM3D [[github.com/geodynamics/specfem3d](https://github.com/geodynamics/specfem3d)] *Spectral Finite Element Method in 3D*, a modeling tool for acoustic, elastic, coupled acoustic/elastic, poroelastic, or seismic wave propagation in any type of conforming mesh of hexahedra, structure or unstructured
- CLAWPACK [[clawpack.org](http://clawpack.org)] *Conservation Laws Package*, a toolkit for linear and nonlinear hyperbolic systems of conservation laws, implementing high-resolution Godunov type methods using limiters
- PyCLAW [[clawpack.org/pyclaw/](http://clawpack.org/pyclaw/)] a Python-based interface to the algorithms of Clawpack and SharpClaw, and the PetClaw package, which adds parallelism through PETSc

# ECRC-curated software

- `petsc4py` [[pypi.org/project/petsc4py/](https://pypi.org/project/petsc4py/)] Python bindings for PETSc
- `slepc4py` [[pypi.org/project/slepc4py/](https://pypi.org/project/slepc4py/)] Python bindings for the eigensolver package SLEPc, based on PETSc solvers
- `mpi4py` [[pypi.org/project/mpi4py/](https://pypi.org/project/mpi4py/)] Python bindings for the Message Passing Interface standard
- MFEM [[mfem.org](https://mfem.org)] a lightweight, scalable C++ library for Finite Element Methods, supported out of Lawrence Livermore National Laboratory, a workhorse code of the US DOE, which uses PETSc solvers
- OpenFOAM [[openfoam.com](https://openfoam.com)] Open Source Computational Fluid Dynamics Toolbox, modeling software for complex fluid flows involving chemical reactions, turbulence and heat transfer, acoustics, solid mechanics and electromagnetics, which uses PETSc solvers

# ECRC software not under curation

- KFUN3D [[github.com/ecrc/kfun3d](https://github.com/ecrc/kfun3d)] an adaption of the NASA code FUN3D for unstructured CFD to manycore architectures, exploiting AVX-512 instructions, which uses PETSc solvers
- BEMFMM [[github.com/ecrc/bemfmm](https://github.com/ecrc/bemfmm)] an FMM-accelerated boundary integral equation solver for wave scattering
- MOAO [[github.com/ecrc/moao](https://github.com/ecrc/moao)] *Multi-Objective Adaptive Optics*, a simulation framework for controlling adaptive optics in earth-based telescopes for manycore architectures and GPUs
- ACR [[github.org/ecrc/acr](https://github.org/ecrc/acr)] *Accelerated Cyclic Reduction*, a cyclic reduction solver for structured linear systems that exploits low rank compression of Schur complements
- RDF-EXP [[github.com/ecrc/rdf-exp](https://github.com/ecrc/rdf-exp)] a test harness for distributed SPARQL engines for Resource Descriptor Framework (RDF) database processing

# ECRC software not under curation

- iCentral [[github.com/fjamour/icentral](https://github.com/fjamour/icentral)] software for detecting and updating “betweenness centrality” in evolving large graphs
- GraMi [[github.com/ehab-abdelhamid/GraMi](https://github.com/ehab-abdelhamid/GraMi)] a framework for frequent subgraph mining in large graphs
- ScaleMine [<https://github.com/ehab-abdelhamid/ScaleMine>] a scalable parallel version of frequent subgraph mining
- AdPart [[github.com/ibrahimabdelaziz/AdPart](https://github.com/ibrahimabdelaziz/AdPart)] *Adaptive Partitioning*, a framework for accelerating SPARQL queries by exploiting hash-based locality

# ECRC software featured at SC'21

**ExaGeoStat**  
A HIGH PERFORMANCE MULTIOBJECT ADAPTIVE OPTICS FRAMEWORK FOR GROUND-BASED ASTRONOMY

ExaGeoStat is a high performance software framework for computational geostatistics in massive systems. The project aims at solving the distributed problem for a given spatial data by parallelizing the computation of the covariance matrix and the estimation of the posterior distribution. ExaGeoStat is designed to be scalable and efficient, and it is implemented in C++ and OpenMP. The software is available on GitHub: <https://github.com/ECRC/ExaGeoStat>

**MOAO**  
A HIGH PERFORMANCE MULTIOBJECT ADAPTIVE OPTICS FRAMEWORK FOR GROUND-BASED ASTRONOMY

MOAO is a high performance software framework for computational geostatistics in massive systems. The project aims at solving the distributed problem for a given spatial data by parallelizing the computation of the covariance matrix and the estimation of the posterior distribution. MOAO is designed to be scalable and efficient, and it is implemented in C++ and OpenMP. The software is available on GitHub: <https://github.com/ECRC/MOAO>

in NVIDIA cuBLAS

**KBLAS**  
A HIGH PERFORMANCE MULTIOBJECT ADAPTIVE OPTICS FRAMEWORK FOR GROUND-BASED ASTRONOMY

KBLAS is a high performance software framework for computational geostatistics in massive systems. The project aims at solving the distributed problem for a given spatial data by parallelizing the computation of the covariance matrix and the estimation of the posterior distribution. KBLAS is designed to be scalable and efficient, and it is implemented in C++ and OpenMP. The software is available on GitHub: <https://github.com/ECRC/KBLAS>

in Cray LibSci

**KSVD**  
A HIGH PERFORMANCE MULTIOBJECT ADAPTIVE OPTICS FRAMEWORK FOR GROUND-BASED ASTRONOMY

KSVD is a high performance software framework for computational geostatistics in massive systems. The project aims at solving the distributed problem for a given spatial data by parallelizing the computation of the covariance matrix and the estimation of the posterior distribution. KSVD is designed to be scalable and efficient, and it is implemented in C++ and OpenMP. The software is available on GitHub: <https://github.com/ECRC/KSVD>

in ExaWave

**GIRIH**  
A HIGH PERFORMANCE MULTIOBJECT ADAPTIVE OPTICS FRAMEWORK FOR GROUND-BASED ASTRONOMY

GIRIH is a high performance software framework for computational geostatistics in massive systems. The project aims at solving the distributed problem for a given spatial data by parallelizing the computation of the covariance matrix and the estimation of the posterior distribution. GIRIH is designed to be scalable and efficient, and it is implemented in C++ and OpenMP. The software is available on GitHub: <https://github.com/ECRC/GIRIH>

in PETSc

**H2Opus**  
A HIGH PERFORMANCE MULTIOBJECT ADAPTIVE OPTICS FRAMEWORK FOR GROUND-BASED ASTRONOMY

H2Opus is a high performance software framework for computational geostatistics in massive systems. The project aims at solving the distributed problem for a given spatial data by parallelizing the computation of the covariance matrix and the estimation of the posterior distribution. H2Opus is designed to be scalable and efficient, and it is implemented in C++ and OpenMP. The software is available on GitHub: <https://github.com/ECRC/H2Opus>

**STARS-H**  
A HIGH PERFORMANCE MULTIOBJECT ADAPTIVE OPTICS FRAMEWORK FOR GROUND-BASED ASTRONOMY

STARS-H is a high performance software framework for computational geostatistics in massive systems. The project aims at solving the distributed problem for a given spatial data by parallelizing the computation of the covariance matrix and the estimation of the posterior distribution. STARS-H is designed to be scalable and efficient, and it is implemented in C++ and OpenMP. The software is available on GitHub: <https://github.com/ECRC/STARS-H>

**AL4SAN**  
A HIGH PERFORMANCE MULTIOBJECT ADAPTIVE OPTICS FRAMEWORK FOR GROUND-BASED ASTRONOMY

AL4SAN is a high performance software framework for computational geostatistics in massive systems. The project aims at solving the distributed problem for a given spatial data by parallelizing the computation of the covariance matrix and the estimation of the posterior distribution. AL4SAN is designed to be scalable and efficient, and it is implemented in C++ and OpenMP. The software is available on GitHub: <https://github.com/ECRC/AL4SAN>

**HiCMA**  
A HIGH PERFORMANCE MULTIOBJECT ADAPTIVE OPTICS FRAMEWORK FOR GROUND-BASED ASTRONOMY

HiCMA is a high performance software framework for computational geostatistics in massive systems. The project aims at solving the distributed problem for a given spatial data by parallelizing the computation of the covariance matrix and the estimation of the posterior distribution. HiCMA is designed to be scalable and efficient, and it is implemented in C++ and OpenMP. The software is available on GitHub: <https://github.com/ECRC/HiCMA>

**MLBS**  
A HIGH PERFORMANCE MULTIOBJECT ADAPTIVE OPTICS FRAMEWORK FOR GROUND-BASED ASTRONOMY

MLBS is a high performance software framework for computational geostatistics in massive systems. The project aims at solving the distributed problem for a given spatial data by parallelizing the computation of the covariance matrix and the estimation of the posterior distribution. MLBS is designed to be scalable and efficient, and it is implemented in C++ and OpenMP. The software is available on GitHub: <https://github.com/ECRC/MLBS>

available at <https://github.com/ecrc/>

# ECRC short course in progress

## Theory and Application of Neural Networks

**Professor Jinchao Xu (PennState)** will give a self-contained introduction to the theory of the neural network function class and its application to image classification and numerical solution of partial differential equations (tentatively scheduled for the second week of October).

- \* Definition of the neural network function class as a generalization of classical finite element functions
- \* Deep ReLU neural networks versus the classic piecewise linear finite element functions
- \* Classical approximation theory of neural network functions
- \* New optimal approximate theory of stable neural network functions
- \* Classic machine learning methods: logistic regression and support vector machine
- \* Deep learning: convolutional neural networks (CNN) for image classification
- \* MgNet: a special CNN obtained from a minor modification of the classical multigrid method
- \* Application and error analysis of neural network for numerical solutions of partial differential equations (PDEs)
- \* Numerical quadrature and Rademacher complexity analysis
- \* Old and new training algorithms for machine learning and numerical PDEs

# ECRC short course in progress

## Theory and Application of Neural Networks

### Thumbnail summary:

- The space of functions that map high-dimensional vector inputs to outputs through neural networks via ReLU activation functions is contained in piecewise linear finite elements.
- It can also be enlarged to contain piecewise linear finite elements of the same dimension by growing the neural network sufficiently deep.
- The approximation properties of ReLU are very rich. This is due, in part, to their nonlinearity.
- Though the spaces can be made identical, traditional finite elements and ReLU-based neural networks have very different structure.
- The concepts of a mesh and of a nodal basis that are so fundamental to the efficiency of solution of finite elements are not present, as such, for neural networks.
- Jinchao is pursuing the applicability of optimal solvers, such as multilevel methods, for finite element problems in neural network training.
- So far, no bottom-line breakthroughs in practical training.

# Why the ReLU function

Traditionally, some prevalent non-linear activation functions, like **sigmoid functions** (or logistic) and hyperbolic tangent, are used in neural networks to get activation values corresponding to each **neuron**. Recently, the ReLU function has been used instead to calculate the activation values in traditional neural network or deep neural network paradigms. The reasons of replacing sigmoid and hyperbolic tangent with ReLU consist of:

**Computation saving** - the ReLU function is able to accelerate the training speed of deep neural networks compared to traditional activation functions since the derivative of ReLU is 1 for a positive input. Due to a constant, deep neural networks do not need to take additional time for computing error terms during training phase.

- the ReLU function does not trigger the vanishing gradient problem when the number of layers grows. This is because this function does not have an asymptotic upper and lower bound. Thus, the earliest layer (the first hidden layer) is able to receive the errors coming from the last layers to adjust all weights between layers. By contrast, a traditional activation function like sigmoid is restricted between 0 and 1, so the errors become small for the first hidden layer. This scenario will lead to a poorly trained neural network.

<https://deepai.org/machine-learning-glossary-and-terms/relu>

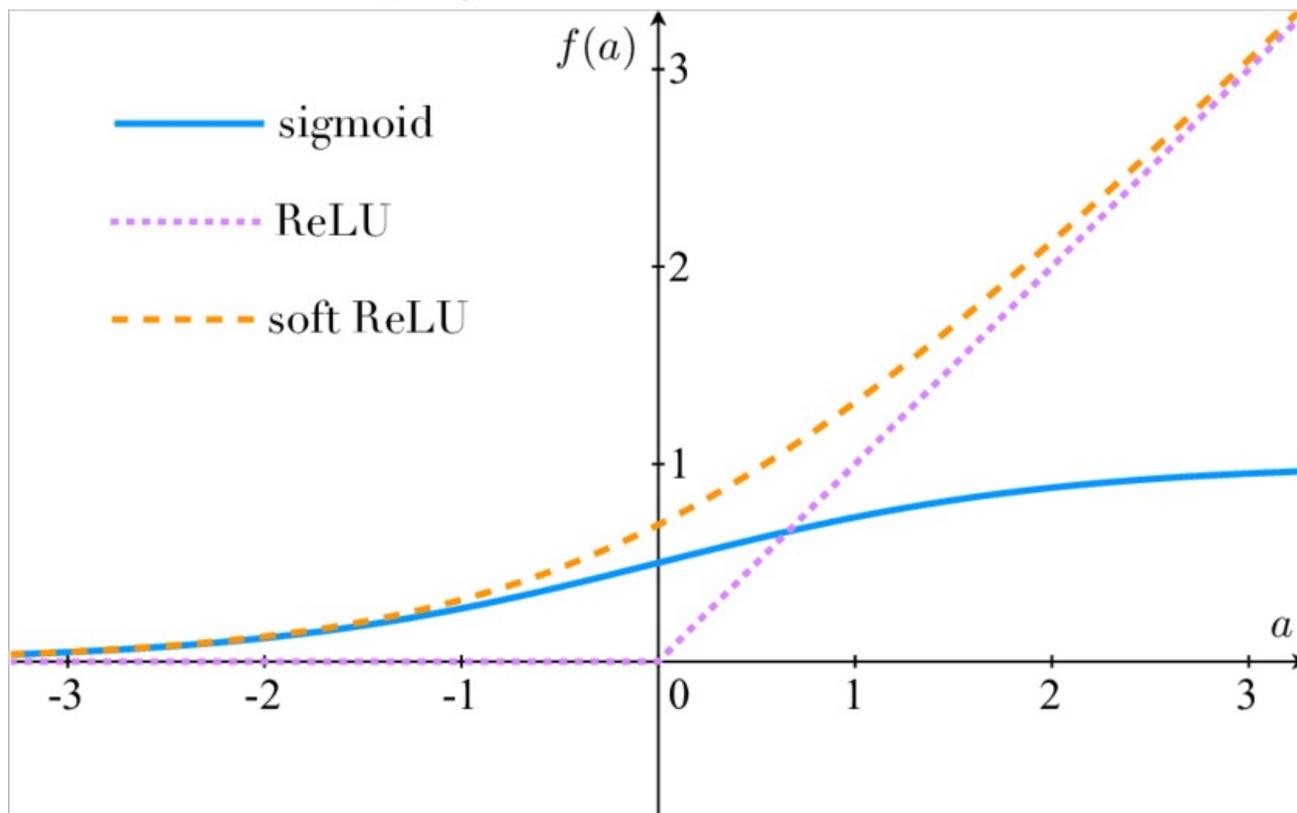
# ReLU (rectified linear unit) function

ReLU is a non-linear **activation function** that is used in multi-layer **neural networks** or deep neural networks. This function can be represented as:

$$f(x) = \max(0, x) \quad (1)$$

where  $x$  = an input value

$$f(x) = \begin{cases} 0, & \text{if } x < 0 \\ x, & \text{if } x \geq 0 \end{cases} \quad (2)$$



# ECRC short course in progress

## Theory and Application of Neural Networks

Not too late to join for the last lecture tomorrow at 9:00am in Lecture Hall 1, Building 9, Room 2322

Topics:

A departure from the theorem-proof style of the first two lectures to look at image classification and PDEs examples.

Slides will be available.

The lectures have been recorded.

# Majlis

