



RESEARCH CENTER
Bordeaux - Sud-Ouest

FIELD

Activity Report 2016

Section Software

Edition: 2017-08-25

1. CAGIRE Project-Team	4
2. CARDAMOM Project-Team	7
3. CARMEN Project-Team	15
4. CQFD Project-Team	19
5. FLOWERS Project-Team	21
6. GEOSTAT Project-Team	34
7. HIEPACS Project-Team	35
8. LFANT Project-Team	39
9. MAGIQUE-3D Project-Team	43
10. MANAO Project-Team	45
11. MEMPHIS Project-Team	47
12. MNEMOSYNE Project-Team	49
13. MONC Project-Team	52
14. PHOENIX Project-Team	54
15. PLEIADE Team	60
16. POSET Team	63
17. POTIOC Project-Team	64
18. REALOPT Project-Team	65
19. SISTM Project-Team	66
20. STORM Team	68
21. TADAAM Team	72

CAGIRE Project-Team

6. New Software and Platforms

6.1. AeroSol

Participants: Simon Delmas [Université de Bordeaux], Sébastien de Brye [Université de Bordeaux], Benjamin Lux [Cagire], Nikolaos Pattakos [Cardamom], Vincent Perrier [Cagire, correspondent], Mario Ricchiuto [Cardamom].

Developed since 2011 by V. Perrier in partnership with the Cardamom Inria team, the AeroSol library is a high order finite element library written in C++. The code design has been carried for being able to perform efficient computations, with continuous and discontinuous finite element methods on hybrid and possibly curvilinear meshes. The work of the Cardamom team is focused on continuous finite element methods, while we focus on discontinuous Galerkin methods. However, everything is done for sharing the largest possible part of code. The distribution of the unknowns is made with the software PaMPA, first developed within the Inria teams Bacchus and Castor, and currently maintained in the Tadaam team.

The generic features of the library are Adaptive wall treatment for a second moment closure in the industrial context

- **High order.** It can be theoretically any order of accuracy, but the finite element basis, and quadrature formula are implemented for having up to a fifth order of accuracy.
- **Hybrid and curvilinear meshes.** AeroSol can deal with up to fifth order conformal meshes composed of lines, triangles, quadrangles, tetrahedra, hexahedra, prism, and pyramids.
- **Continuous and discontinuous discretization.** AeroSol deals with both continuous and discontinuous finite element methods.

We would like to emphasize three assets of this library:

- **Its development environment** For allowing a good collaborative work and a functional library, a strong emphasis has been put on the use of modern collaborative tools for developing our software. This includes the active use of a repository, the use of CMake for the compilation, the constant development of unitary and functional tests for all the parts of the library (using CTest), and the use of the continuous integration tool Jenkins for testing the different configurations of AeroSol and its dependencies. Efficiency is regularly tested with direct interfacing with the PAPI library or with tools like scalasca.
- **Its genericity** A lot of classes are common to all the discretization, for example classes concerning I/O, finite element functions, quadrature, geometry, time integration, linear solver, models and interface with PaMPA. Adding simple features (e.g. models, numerical flux, finite element basis or quadrature formula) can be easily done by writing the class, and declaring its use in only one class of the code.
- **Its efficiency** This modularity is achieved by means of template abstraction for keeping good performances. Dedicated efficient implementation, based on the data locality of the discontinuous Galerkin method has been developed. As far as parallelism is concerned, we use point-to-point communications, the HDF5 library for parallel I/O. The behavior of the AeroSol library at medium scale (1000 to 2000 cores) was studied in [27].

The AeroSol project fits with the first axis of the Bordeaux Sud Ouest development strategy, which is to build a coherent software suite scalable and efficient on new architectures, as the AeroSol library relies on several tools developed in other Inria teams, especially for the management of the parallel aspects. At the end of 2015, AeroSol had the following features:

- **Boundary conditions** Periodic boundary conditions, time-dependent inlet and outlet boundary conditions. Adiabatic wall and isothermal wall. Steger-Warming based boundary condition. Synthetic Eddy Method for generating turbulence.
- **C++/Fortran interface** Tests for binding fortran with C++.
- **Development environment** An upgraded use of CMake for compilation (gcc, icc and xlc), CTest for automatic tests and memory checking, lcov and gcov for code coverage reports. A CDash server for collecting the unitary tests and the memory checking. An under development interface for functional tests. Optional linking with HDF5, PAPI, with dense small matrices libraries (BLAS, Eigen). An updated shared project Plafrim and joint project Aerosol/Scotch/PaMPA project on the continuous integration platform. An on-going integration of SPack for handling dependencies. A fixed ESSL interface.
- **Finite elements** up to fourth degree for Lagrange finite elements and hierarchical orthogonal finite element basis (with Dubiner transform on simplices) on lines, triangles, quadrangles, tetrahedra, prisms, hexaedra and pyramids. Finite element basis that are interpolation basis on Gauss-Legendre points for lines, quadrangles, and hexaedra, and triangle (only 1st and 2nd order).
- **Geometry** Elementary geometrical functions for first order lines, triangles, quadrangles, prisms, tetrahedra, hexaedra and pyramids. Handling of high order meshes.
- **In/Out** Link with the XML library for handling with parameter files. Parallel reader for GMSH, with an embedded geometrical pre-partitioner. Writer on the VTK-ASCII legacy format (cell and point centered). Parallel output in vtu and pvtu (Paraview) for cell-centered visualization, and XDMF/HDF5 format for both cell and point centered visualization. Ability of saving the high order solution and restarting from it. Computation of volumic and probe statistics. Ability of saving averaged layer data in quad and hexa meshes. Ability of defining user defined output visualization variables.
- **Instrumentation** Aerosol can give some traces on memory consumption/problems with an interfacing with the PAPI library. Tests have also been performed with VTUNE and TAU. Tests with Maqao and Scalasca (VIHPS workshop).
- **Linear Solvers** Link with the external linear solver UMFPack, PETSc and MUMPS. Internal solver for diagonal and block-diagonal matrices.
- **Memory handling** discontinuous and continuous, sequential and parallel discretizations based on PaMPA for generic meshes, including hybrid meshes.
- **Models** Perfect gas Euler system, real gas Euler system (template based abstraction for a generic equation of state), scalar advection, Waves equation in first order formulation, generic interface for defining space-time models from space models. Diffusive models: isotropic and anisotropic diffusion, compressible Navier-Stokes. Scalar advection-diffusion model. Linearized Euler equations, and Sutherland model for non isothermal diffusive flows. Shallow-water model.
- **Multigrid** Development of p -multigrid methods. This includes also the possibility of beginning a computation with an order and to decrease or increase the order of approximation when restarting. For the p multigrid methods, V and W cycle have been developed, and restriction and prolongation operators have also been developed. In progress implementation of h -multigrid, with the development of tests of the aggregation methods of PaMPA, and the definition of finite element basis on arbitrary cells.
- **Numerical fluxes** Centered fluxes, exact Godunov' flux for linear hyperbolic systems, and Lax-Friedrich flux. Riemann solvers for Low Mach flows. Numerical flux accurate for steady and unsteady computations.

- **Numerical schemes** Continuous Galerkin method for the Laplace problem (up to fifth order) with non consistent time iteration or with direct matrix inversion. Explicit and implicit discontinuous Galerkin methods for hyperbolic systems, diffusive and advection-diffusion problems. In progress optimization by stocking the geometry for advection problems. SUPG and Residual distribution schemes. Optimization of DG schemes for advection-diffusion problems: stocking of the geometry and use of BLAS for all the linear phases of the scheme.
- **Parallel computing** Mesh redistribution, computation of Overlap with PaMPA. Collective asynchronous communications (PaMPA based). Asynchronous point to point communications. Tests on the cluster Avakas from MCIA, and on Mésocentre de Marseille, and PlaFRIM. Tier-1 Turing (Blue-Gene). Weighted load balancing for hybrid meshes.
- **Postprocessing** High order projections over line postprocessing, possibility of stocking averaged data, such as the average flow and the Reynolds stresses.
- **Quadrature formula** up to 11th order for Lines, Quadrangles, Hexaedra, Pyramids, Prisms, up to 14th order for tetrahedron, up to 21st order for triangles. Gauss-Lobatto type quadrature formula for lines, triangles, quadrangles and hexaedra.
- **Time iteration** explicit Runge-Kutta up to fourth order, explicit Strong Stability Preserving schemes up to third order. Optimized CFL time schemes: SSP(2,3) and SSP(3,4). CFL time stepping. Implicit integration with BDF schemes from 2nd to 6th order Newton method for stationary problems. Implicit unstationary time iterator non consistent in time for stationary problems. Implementation of in house GMRES and conjugate gradient based on Jacobian free iterations.
- **Validation** Poiseuille, Taylor-Green vortex. Laplace equation on a ring and Poiseuille flow on a ring. Volumic forcing based on wall dissipation. Turbulent channel flow.

In 2016, the following features have been added:

- Geometric multigrid methods: aggregation of the mesh based on PaMPA, definition of finite element basis on arbitrary shape cells. Definition of geometry, quadratures and numerical schemes on aggregated finite elements.
- Sutherland law in the Navier-Stokes equations.
- Mass matrix free implementation of discontinuous Galerkin methods.
- Improvement of installation documentation. Spack based installation.
- Implementation of Boussinesq type models and Shallow water discretizations with well balancing, positivity preserving, wet-dry handling, limiters based on entropy viscosity,
- Implementation of Barotropic Euler equations
- Implementation of Taylor-based basis on simplices.

CARDAMOM Project-Team

5. New Software and Platforms

5.1. AeroSol

Participants: Simon Delmas [Universit de Bordeaux], Sbastien de Brye [Universit de Bordeaux], Benjamin Lux [Cagire], Nikolaos Pattakos [Cardamom], Vincent Perrier [Cagire, correspondent], Mario Ricchiuto [Cardamom].

Developed since 2011 by V. Perrier in partnership with the Cardamom Inria team, the AeroSol library is a high order finite element library written in C++. The code design has been carried for being able to perform efficient computations, with continuous and discontinuous finite element methods on hybrid and possibly curvilinear meshes.

The work of the Cardamom team is focused on continuous finite element methods, while we focus on discontinuous Galerkin methods. However, everything is done for sharing the largest possible part of code. The distribution of the unknowns is made with the software PaMPA, first developed within the Inria teams Bacchus and Castor, and currently maintained in the Tadaam team.

The generic features of the library are Adaptive wall treatment for a second moment closure in the industrial context

- **High order.** It can be theoretically any order of accuracy, but the finite element basis, and quadrature formula are implemented for having up to a fifth order of accuracy.
- **Hybrid and curvilinear meshes.** AeroSol can deal with up to fifth order conformal meshes composed of lines, triangles, quadrangles, tetrahedra, hexahedra, prism, and pyramids.
- **Continuous and discontinuous discretization.** AeroSol deals with both continuous and discontinuous finite element methods.

We would like to emphasize three assets of this library:

- **Its development environment** For allowing a good collaborative work and a functional library, a strong emphasis has been put on the use of modern collaborative tools for developing our software. This includes the active use of a repository, the use of CMake for the compilation, the constant development of unitary and functional tests for all the parts of the library (using CTest), and the use of the continuous integration tool Jenkins for testing the different configurations of AeroSol and its dependencies. Efficiency is regularly tested with direct interfacing with the PAPI library or with tools like scalasca.
- **Its genericity** A lot of classes are common to all the discretization, for example classes concerning I/O, finite element functions, quadrature, geometry, time integration, linear solver, models and interface with PaMPA. Adding simple features (e.g. models, numerical flux, finite element basis or quadrature formula) can be easily done by writing the class, and declaring its use in only one class of the code.
- **Its efficiency** This modularity is achieved by means of template abstraction for keeping good performances. Dedicated efficient implementation, based on the data locality of the discontinuous Galerkin method has been developed. As far as parallelism is concerned, we use point-to-point communications, the HDF5 library for parallel I/O. The behavior of the AeroSol library at medium scale (1000 to 2000 cores) was studied in [108].

The AeroSol project fits with the first axis of the Bordeaux Sud Ouest development strategy, which is to build a coherent software suite scalable and efficient on new architectures, as the AeroSol library relies on several tools developed in other Inria teams, especially for the management of the parallel aspects. At the end of 2015, AeroSol had the following features:

- **Boundary conditions** Periodic boundary conditions, time-dependent inlet and outlet boundary conditions. Adiabatic wall and isothermal wall. Steger-Warming based boundary condition. Synthetic Eddy Method for generating turbulence.
- **C++/Fortran interface** Tests for binding fortran with C++.
- **Development environment** An upgraded use of CMake for compilation (gcc, icc and xlc), CTest for automatic tests and memory checking, lcov and gcov for code coverage reports. A CDash server for collecting the unitary tests and the memory checking. An under development interface for functional tests. Optional linking with HDF5, PAPI, with dense small matrices libraries (BLAS, Eigen). An updated shared project Plafrim and joint project Aerosol/Scotch/PaMPA project on the continuous integration platform. An on-going integration of SPack for handling dependencies. A fixed ESSL interface.
- **Finite elements** up to fourth degree for Lagrange finite elements and hierarchical orthogonal finite element basis (with Dubiner transform on simplices) on lines, triangles, quadrangles, tetrahedra, prisms, hexaedra and pyramids. Finite element basis that are interpolation basis on Gauss-Legendre points for lines, quadrangles, and hexaedra, and triangle (only 1st and 2nd order).
- **Geometry** Elementary geometrical functions for first order lines, triangles, quadrangles, prisms, tetrahedra, hexaedra and pyramids. Handling of high order meshes.
- **In/Out** Link with the XML library for handling with parameter files. Parallel reader for GMSH, with an embedded geometrical pre-partitioner. Writer on the VTK-ASCII legacy format (cell and point centered). Parallel output in vtu and pvtu (Paraview) for cell-centered visualization, and XDMF/HDF5 format for both cell and point centered visualization. Ability of saving the high order solution and restarting from it. Computation of volumic and probe statistics. Ability of saving averaged layer data in quad and hexa meshes. Ability of defining user defined output visualization variables.
- **Instrumentation** Aerosol can give some traces on memory consumption/problems with an interfacing with the PAPI library. Tests have also been performed with VTUNE and TAU. Tests with Maqao and Scalasca (VIHPS workshop).
- **Linear Solvers** Link with the external linear solver UMFPack, PETSc and MUMPS. Internal solver for diagonal and block-diagonal matrices.
- **Memory handling** discontinuous and continuous, sequential and parallel discretizations based on PaMPA for generic meshes, including hybrid meshes.
- **Models** Perfect gas Euler system, real gas Euler system (template based abstraction for a generic equation of state), scalar advection, Waves equation in first order formulation, generic interface for defining space-time models from space models. Diffusive models: isotropic and anisotropic diffusion, compressible Navier-Stokes. Scalar advection-diffusion model. Linearized Euler equations, and Sutherland model for non isothermal diffusive flows. Shallow-water model.
- **Multigrid** Development of p -multigrid methods. This includes also the possibility of beginning a computation with an order and to decrease or increase the order of approximation when restarting. For the p multigrid methods, V and W cycle have been developed, and restriction and prolongation operators have also been developed. In progress implementation of h -multigrid, with the development of tests of the aggregation methods of PaMPA, and the definition of finite element basis on arbitrary cells.
- **Numerical fluxes** Centered fluxes, exact Godunov' flux for linear hyperbolic systems, and Lax-Friedrich flux. Riemann solvers for Low Mach flows. Numerical flux accurate for steady and unsteady computations.

- **Numerical schemes** Continuous Galerkin method for the Laplace problem (up to fifth order) with non consistent time iteration or with direct matrix inversion. Explicit and implicit discontinuous Galerkin methods for hyperbolic systems, diffusive and advection-diffusion problems. In progress optimization by stocking the geometry for advection problems. SUPG and Residual distribution schemes. Optimization of DG schemes for advection-diffusion problems: stocking of the geometry and use of BLAS for all the linear phases of the scheme.
- **Parallel computing** Mesh redistribution, computation of Overlap with PaMPA. Collective asynchronous communications (PaMPA based). Asynchronous point to point communications. Tests on the cluster Avakas from MCIA, and on Mésocentre de Marseille, and PlaFRIM. Tier-1 Turing (Blue-Gene). Weighted load balancing for hybrid meshes.
- **Postprocessing** High order projections over line postprocessing, possibility of stocking averaged data, such as the average flow and the Reynolds stresses.
- **Quadrature formula** up to 11th order for Lines, Quadrangles, Hexaedra, Pyramids, Prisms, up to 14th order for tetrahedron, up to 21st order for triangles. Gauss-Lobatto type quadrature formula for lines, triangles, quadrangles and hexaedra.
- **Time iteration** explicit Runge-Kutta up to fourth order, explicit Strong Stability Preserving schemes up to third order. Optimized CFL time schemes: SSP(2,3) and SSP(3,4). CFL time stepping. Implicit integration with BDF schemes from 2nd to 6th order Newton method for stationary problems. Implicit unstationary time iterator non consistent in time for stationary problems. Implementation of in house GMRES and conjugate gradient based on Jacobian free iterations.
- **Validation** Poiseuille, Taylor-Green vortex. Laplace equation on a ring and Poiseuille flow on a ring. Volumic forcing based on wall dissipation. Turbulent channel flow.

In 2016, the following features have been added:

- Geometric multigrid methods: aggregation of the mesh based on PaMPA, definition of finite element basis on arbitrary shape cells. Definition of geometry, quadratures and numerical schemes on aggregated finite elements.
- Sutherland law in the Navier-Stokes equations.
- Mass matrix free implementation of discontinuous Galerkin methods.
- Improvement of installation documentation. Spack based installation.
- Implementation of Boussinesq type models and Shallow water discretizations with well balancing, positivity preserving, wet-dry handling, limiters based on entropy viscosity,
- Implementation of Barotropic Euler equations
- Implementation of Taylor-based basis on simplices.

5.2. CAKE

KEYWORDS: Moving bodies, Rarefied flows

- Contact: Luc Mieussens
- URL: <http://www.cake-solver.com/>

FUNCTIONAL DESCRIPTION

Cake (Cut cell Algorithm for Kinetic Equations) can simulate 2D plane rarefied flows around moving obstacles, by using an immersed boundary technique with Cartesian grids. This code can simulate flows induced by temperature gradients, like the thermal creep flow. It has for instance been applied to the simulation of the Crookes radiometer.

5.3. Crysa

KEYWORDS: Image processing, structural analysis, 2D crystallography

- Participants: Jean Mercat and Cécile Dobrzynski
- Partners: ISM - LCPO - LCTS (UMR 5801)
- Contact: Cécile Dobrzynski

Crysa is a library allowing to study the organization of objects once placed in an hexagonal grid thus allowing to analyze the crystal structure/organization of an image. The library allows to detect regions of coherence in an image of crystals, and it allows to determine e.g. the good separation of objects in an experiment (new polymeric materials, plasma, ...).

5.4. Cut-ANOVA

KEYWORDS: Stochastic models - Uncertainty quantification

- Participants: Pietro-Marco Congedo
- Contact: Pietro-Marco Congedo

SCIENTIFIC DESCRIPTION

An anchored analysis of variance (ANOVA) method is proposed to decompose the statistical moments. Compared to the standard ANOVA with mutually orthogonal component functions, the anchored ANOVA, with an arbitrary choice of the anchor point, loses the orthogonality if employing the same measure. However, an advantage of the anchored ANOVA consists in the considerably reduced number of deterministic solver's computations, which renders the uncertainty quantification of real engineering problems much easier. Different from existing methods, the covariance decomposition of the output variance is used in this work to take account of the interactions between non-orthogonal components, yielding an exact variance expansion and thus, with a suitable numerical integration method, provides a strategy that converges. This convergence is verified by studying academic tests. In particular, the sensitivity problem of existing methods to the choice of anchor point is analyzed via the Ishigami case, and we point out that covariance decomposition survives from this issue. Also, with a truncated anchored ANOVA expansion, numerical results prove that the proposed approach is less sensitive to the anchor point. The covariance-based sensitivity indices (SI) are also used, compared to the variance-based SI. Furthermore, we emphasize that the covariance decomposition can be generalized in a straightforward way to decompose higher-order moments. For academic problems, results show the method converges to exact solution regarding both the skewness and kurtosis. The proposed method can indeed be applied to a large number of engineering problems.

FUNCTIONAL DESCRIPTION

The Cut-ANOVA code (Fortran 90, MPI + OpenMP) is devoted to the stochastic analysis of numerical simulations. The method implemented is based on the spectral expansion of "anchored ANOVA", allowing the covariance-based sensitivity analysis. Compared to the conventional Sobol method, "Cut-ANOVA" provides three sensitivity indices instead of one, which allows a better analysis of the reliability of the numerical prediction. On the other hand, "Cut-ANOVA" is able to compute the higher order statistical moments such as the Skewness (3-rd order moment) and Kurtosis (4-th order moment). Several dimension reduction techniques have also been implemented to reduce the computational cost. Finally, thanks to the innovative method implemented into the Code Cut-ANOVA, one can obtain a similar accuracy for stochastic quantities by using a considerably less number of deterministic model evaluations, compared with the classical Monte Carlo method.

5.5. Fmg

KEYWORDS: Mesh adaptation, mesh deformation, elasticity, laplacian mesh equation

- Participants: Cécile Dobrzynski, Mario Ricchiuto, Leo Nouveau and Luca Arpaia
- Contact: Cécile Dobrzynski

FUNCTIONAL DESCRIPTION

FMG is a library deforming an input/reference simplicial mesh w.r.t. a given smoothness error monitor (function gradient or Hessian), metric field, or given mesh size distribution. Displacements are computed by solving an elliptic Laplacian type equation with a continuous finite element method. The library returns an adapted mesh with a corresponding projected solution, obtained by either a second order projection, or by an ALE finite element remap. This year a new semi-linear elasticity formulation has been implemented involving a constant coefficient PDE with a nonlinear *force* accounting for the smoothness of the target function. The advantage of this approach is that the non-linearity does not influence the elastic differential operator thus leading to a description of the deformation governed by a time-invariant matrix even in unsteady simulations. Other developments currently implemented in SLOWS are being imported in FMG.

5.6. MMG3D

KEYWORDS: Mesh - Anisotropic - Mesh adaptation

- Participants: Cécile Dobrzynski, Pascal Frey, Charles Dapogny and Algiane Froehly
- Partners: CNRS - IPB - Université de Bordeaux - UPMC
- Contact: Cécile Dobrzynski
- URL: <http://www.mmgtools.org>

SCIENTIFIC DESCRIPTION

Mmg3d is an open source software for tetrahedral remeshing. It performs local mesh modifications. The mesh is iteratively modified until the user prescriptions satisfaction.

Mmg3d can be used by command line or using the library version (C, C++ and Fortran API) : - It is a new version of the MMG3D4 software. It remesh both the volume and surface mesh of a tetrahedral mesh. It performs isotropic and anisotropic mesh adaptation and isovalue discretization of a level-set function.

Mmg3d allows to control the boundaries approximation: The "ideal" geometry is reconstructed from the piecewise linear mesh using cubic Bezier triangular patches. The surface mesh is modified to respect a maximal Hausdorff distance between the ideal geometry and the mesh.

Inside the volume, the software performs local mesh modifications (such as edge swap, pattern split, isotropic and anisotropic Delaunay insertion...).

FUNCTIONAL DESCRIPTION

Mmg3d is one of the software of the Mmg platform. It is dedicated to the modification of 3D volume meshes. It performs the adaptation and the optimization of a tetrahedral mesh and allows to discretize an isovalue.

Mmg3d performs local mesh modifications. The mesh is iteratively modified until the user prescription satisfaction.

5.7. Mmg platform

KEYWORDS: Mesh - Mesh generation - Anisotropic - Mesh adaptation - Isovalue discretization

- Participants: Cécile Dobrzynski, Charles Dapogny, Pascal Frey and Algiane Froehly
- Partners: CNRS - IPB - Université de Bordeaux - UPMC
- Contact: Cécile Dobrzynski
- URL: <http://www.mmgtools.org>

SCIENTIFIC DESCRIPTION

The Mmg platform gathers open source software for two-dimensional, surface and volume remeshing. The platform software performs local mesh modifications. The mesh is iteratively modified until the user prescriptions satisfaction.

The 3 softwares can be used by command line or using the library version (C, C++ and Fortran API) : - Mmg2d performs mesh generation and isotropic and anisotropic mesh adaptation. - Mmgs allows isotropic and anisotropic mesh adaptation for 3D surface meshes. - Mmg3d is a new version of the MMG3D4 software. It remesh both the volume and surface mesh of a tetrahedral mesh. It performs isotropic and anisotropic mesh adaptation and isovalue discretization of a level-set function.

The platform software allow to control the boundaries approximation: The "ideal" geometry is reconstruct from the piecewise linear mesh using cubic Bezier triangular patches. The surface mesh is modified to respect a maximal Hausdorff distance between the ideal geometry and the mesh.

Inside the volume, the software perform local mesh modifications (such as edge swap, pattern split, isotropic and anisotropic Delaunay insertion...).

FUNCTIONAL DESCRIPTION

The Mmg platform gathers open source software for two-dimensional, surface and volume remeshing. It provides three applications : 1) mmg2d: generation of a triangular mesh , adaptation and optimization of a triangular mesh 2) mmgs: adaptation and optimization of a surface triangulation representing a piecewise linear approximation of an underlying surface geometry 3) mmg3d: adaptation and optimization of a tetrahedral mesh and isovalue discretization

The platform software perform local mesh modifications. The mesh is iteratively modified until the user prescription satisfaction.

5.8. NOMESH

- Participants: Cécile Dobrzynski and Algiane Froehly
- Partners: CNRS - IPB - Université de Bordeaux
- Contact: Cécile Dobrzynski

KEYWORDS: Mesh - Curved mesh - Tetrahedral mesh

FUNCTIONAL DESCRIPTION

NOMESH is a software allowing the generation of three order curved simplicial meshes. Starting from a "classical" mesh with straight elements composed by triangles and/or tetrahedra, we are able to curve the boundary mesh. Starting from a mesh with some curved elements, we can verify if the mesh is valid, that means there is no crossing elements and only positive Jacobian. If the curved mesh is non valid, we modify it using linear elasticity equations until having a valid curved mesh.

5.9. SLOWS

KEYWORDS: Free surface flows, Unstructured meshes, shallow water equations, Boussinesq equations

- Participants: Luca Arpaia, Andrea Filippini, Maria Kazolea, Mario Ricchiuto and Nikolaos Pattakos
- Contact: Mario Ricchiuto

FUNCTIONAL DESCRIPTION

SLOWS (Shallow-water fLOWS) is a C-platform allowing the simulation of free surface shallow water flows with friction. It can be used to simulate near shore hydrodynamics, wave transformations processes, etc. The kernel of the CODE is a shallow water solver based on second order residual distribution or second and third order finite volume schemes. Three different approaches are available to march in time, based on conditionally depth-positivity preserving implicit schemes, or on conditionally depth-positivity preserving genuinely explicit discretizations, or on an unconditionally depth-positivity preserving space-time approach. Newton and frozen Newton loops are used to solve the implicit nonlinear equations. Linear algebraic sparse systems arising in the discretization are solved with This year several enhancements have been implemented

- a correction of both the residual distribution and finite volume methods to solve the shallow water equations in spherical (or Mercator) curvilinear coordinates;

- mass-conserving mesh movement to adapt an initial grid to wet-dry interfaces as well as to other physical features of the flow;
- a new library has been developed to enhance the shallow water equations. This library computes an algebraic source term by inverting an elliptic (grad-div type) PDE. The addition of this term to the shallow water version of SLOWS allows to recover the a fully nonlinear weakly dispersive Green-Naghdi solver. The solution of the elliptic PDE is performed with a classical Galerkin FEM approach, and MUMPS is the MUMPS library is used to invert the resulting matrix;
- Initial optimization and OpenMP parallelization of the shallow water kernel.

SLOWS is our main simulation tool in both the TANDEM and Tides projects.

5.10. Sparse-PDD

Adaptive sparse polynomial dimensional decomposition for global sensitivity analysis

KEYWORDS: Stochastic models - Uncertainty quantification

- Participants: Pietro-Marco Congedo
- Contact: Pietro-Marco Congedo

SCIENTIFIC DESCRIPTION

The polynomial dimensional decomposition (PDD) is employed in this code for the global sensitivity analysis and uncertainty quantification (UQ) of stochastic systems subject to a moderate to large number of input random variables. Due to the intimate structure between the PDD and the Analysis of Variance (ANOVA) approach, PDD is able to provide a simpler and more direct evaluation of the Sobol sensitivity indices, when compared to the Polynomial Chaos expansion (PC). Unfortunately, the number of PDD terms grows exponentially with respect to the size of the input random vector, which makes the computational cost of standard methods unaffordable for real engineering applications. In order to address the problem of the curse of dimensionality, this code proposes essentially variance-based adaptive strategies aiming to build a cheap meta-model (i.e. surrogate model) by employing the sparse PDD approach with its coefficients computed by regression. Three levels of adaptivity are carried out in this code: 1) the truncated dimensionality for ANOVA component functions, 2) the active dimension technique especially for second- and higher-order parameter interactions, and 3) the stepwise regression approach designed to retain only the most influential polynomials in the PDD expansion. During this adaptive procedure featuring stepwise regressions, the surrogate model representation keeps containing few terms, so that the cost to resolve repeatedly the linear systems of the least-square regression problem is negligible. The size of the finally obtained sparse PDD representation is much smaller than the one of the full expansion, since only significant terms are eventually retained. Consequently, a much less number of calls to the deterministic model is required to compute the final PDD coefficients.

FUNCTIONAL DESCRIPTION

This code allows an efficient meta-modeling for a complex numerical system featuring a moderate-to-large number of uncertain parameters. This innovative approach involves polynomial representations combined with the Analysis of Variance decomposition, with the objective to quantify the numerical output uncertainty and its sensitivity upon the variability of input parameters.

5.11. TUCWave

KEYWORDS: free surface flows, Boussinesq equations, weakly nonlinear models, unstructured grids

- Participants: Maria Kazolea, Argiris Delis and Ioannis Nikolos
- Contact: Maria Kazolea

FUNCTIONAL DESCRIPTION

Fortran Planform which accounts for the study of near shore processes. TUCWave uses a high-order well-balanced unstructured finite volume (FV) scheme on triangular meshes for modeling weakly nonlinear and weakly dispersive water waves over varying bathymetries, as described by the 2D depth-integrated extended Boussinesq equations of Nwogu (1993), rewritten in conservation law form. The FV scheme numerically solves the conservative form of the equations following the median dual node-centered approach, for both the advective and dispersive part of the equations. The code developed follows an efficient edge based structured technique. For the advective fluxes, the scheme utilizes an approximate Riemann solver along with a well-balanced topography source term up-winding. Higher order accuracy in space and time is achieved through a MUSCL-type reconstruction technique and through a strong stability preserving explicit Runge-Kutta time stepping. Special attention is given to the accurate numerical treatment of moving wet/dry fronts and boundary conditions. Furthermore, the model is applied to several examples of wave propagation over variable topographies and the computed solutions are compared to experimental data. TUCWave is used in the TANDEM project to provide reference solutions with a weakly nonlinear and dispersive model.

CARMEN Project-Team

6. New Software and Platforms

6.1. CEPS: a Cardiac ElectroPhysiology Simulator

The Carmen team develops a software library to perform high-performance numerical simulations in cardiac electrophysiology using unstructured three-dimensional grids. The software, called CEPS (*Cardiac Electro-physiology Simulation*) is developed as a common tool for researchers in the Carmen team and for our partners and colleagues in scientific computing and biomedical engineering. The goal of CEPS is to facilitate the development of new numerical methods and new physical models.

Compared to other existing software, CEPS aims at providing a more general framework of integration for new methods or models and a better efficiency in parallel. CEPS is designed to run on massively parallel architectures, and to make use of state-of-the-art and well-known computing libraries to achieve realistic and complex heart simulations. The largest part of CEPS was developed by the Junior Engineer M. Juhoor, supervised by N. Zemzemi, during the CEPS ADT (*Action de Développement Technologique*).

To enforce a sound development process, some engineering and validation tools are used:

- Git hosted at the Inria GForge ([ceps](#)) to manage versions;
- Cmake for the building process
- [Jenkins](#), hosted at the Inria continuous integration service, which runs a test suite of about 200 tests after every commit.

Main users and developers of CEPS are the PhD students of Carmen, i.e.

- A. Gérard, who uses CEPS for patient-specific modeling, has implemented a bilayer model using coupled nodes.
- Charlie Douanla-Lountsi currently works on high-order temporal integration methods, for later integration in CEPS.
- P. E. Bécue is developing a code to run microscopic-scale models (section 3.4) and wrote a coupled node assembler to support this work.

Since January 2015, M. Fuentes from the *Service d'Experimentation et de Développement* (SED), is responsible for developing new features in CEPS, improve robustness, efficiency, and documentation. M. Juhoor, who has previously worked on CEPS, and works on the IDAM project, brings us from time to time his expertise. Actions done in 2016 include:

- support for P2 Lagrange finite Elements
- node partitioning using the PTScotch partitioner
- input files and VtkReader (M. Juhoor)
- code refactoring
- documentation writing

6.2. IDAM

The goal of the IDAM project is to define a collection of plugins in the MUSIC software in order to create realistic meshes for the CEPS code. MUSIC is a multimodal platform for cardiac imaging developed by the imaging team at IHU LIRYC (<https://bil.inria.fr/fr/software/view/1885/tab>). Information comes from magnetic resonance imaging and cardiac tomography performed in the clinic and in the LIRYC laboratories. Building complete cardiac models directly from imaging data requires expert knowledge and is time-consuming and error-prone: specific expertise and multiple software tools are often needed to process data stemming from medical imaging into realistic meshes and parameter distributions.

IDAM aims to streamline the workflow of a complete cardiac simulation: anatomical mesh generation from patient-specific data, description of simulation parameters, and eventually analysis of simulation results obtained by simulation packages like CEPS (<https://bil.inria.fr/fr/software/view/2630/tab>). IDAM integrates tools from other Inria teams by using specialized libraries, for example MMG (<https://bil.inria.fr/fr/software/view/2824/tab>) for high-quality mesh generation.




Figure 1. Mesh of a human torso and one of the X-ray computed-tomography slices on which it was based.

6.3. Platforms

6.3.1. Propag-5

Applied modeling studies performed by the Carmen team, especially M. Potse and M. Kania, in collaboration with IHU Liryc and foreign partners [39], [43], [18], [20], [14] [45], rely to a great extent on high-performance computations on the national supercomputers Curie, Occigen, and Turing. While the newly developed CEPS code is not ready to run efficiently on these systems we rely on an older code named Propag-5. This code is the result of a decades-long development first at the *Université de Montréal* in Canada, then at Maastricht University in the Netherlands, and finally at the Institute of Computational Science of the *Università della Svizzera italiana* in Lugano, Switzerland. Relatively small contributions to this code have been made by the Carmen team.

The predecessor of Propag-5, named Propag-4, was developed by M. Potse at the *Université de Montréal* [8]. It was based on earlier model code developed there by the team of Prof. R. Gulrajani [50], [53], and was parallelized with OpenMP to utilize the shared-memory SGI supercomputers available there at the time. Propag-4 was the first code ever able to run a bidomain reaction-diffusion model of the entire human ventricles; a problem 30 times larger than what had been reported before [8].

In order to utilize the more recent distributed-memory architectures Propag-4 was transformed into the hybrid MPI-OpenMP code Propag-5 at the Institute of Computational Science in Lugano by D. Krause and M. Potse [49]. The resulting code has been used for numerous applied studies. An important limitation of the Propag code is that it relies on a semi-structured mesh with a uniform resolution. On the other hand, the code scales excellently to large core counts and, as it is controlled completely with command-line flags and configuration files, it can be used by non-programmers. It also features

- a plugin system for membrane models,
- a completely parallel workflow, including the initial anatomy input and mesh partitioning, which allows it to work with meshes of more than 10^9 nodes,

- a flexible output scheme allowing hundreds of different state variables and transient variables to be output to file, when desired, using any spatial and temporal subsampling,
- a configurable, LUSTRE-aware parallel output system in which groups of processes write HDF5/netCDF files, and
- CWEB documentation of the entire code base.

The code has been stable and reliable for several years, and only minor changes are being made currently. It can be considered the workhorse for our HPC work until CEPS takes over.

6.3.2. Gepetto

Gepetto, named after a famous model maker, is a software suite that transforms a surface mesh of the heart into a set of (semi-)structured meshes for use by the Propag software or others. It creates the different fiber orientations in the model, including the transmurally rotating ventricular fibers and the various bundle structures in the atria (figure 2), and creates layers with possibly different electrophysiological properties across the wall. A practically important function is that it automatically builds the matching heart and torso meshes that Propag uses to simulate potentials in the torso (at a resolution of 1 mm) after projecting simulation results from the heart model (at 0.1 to 0.2 mm) on the coarser torso mesh [52]. Like Propag, the Gepetto software results from a long-term development that started in Montreal, Canada, around 2002. The code for atrial fiber structure was developed by our team.

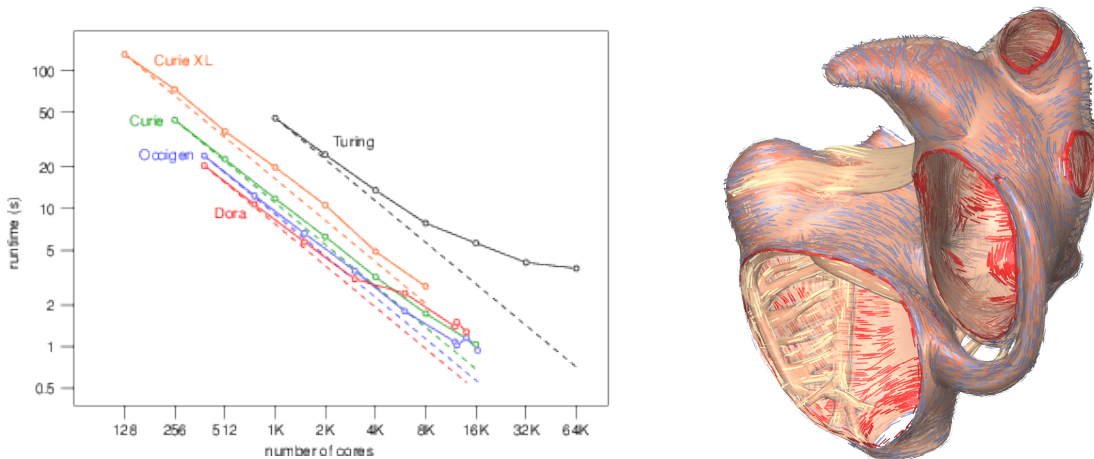


Figure 2. **Left:** Strong scaling of the Propag-5 code on a monodomain reaction-diffusion equation for four systems: The Bull clusters Curie (small nodes and large nodes, XL) and Occigen, The Cray XC30 “Piz Dora” at the Swiss supercomputing center CSCS and the IBM BlueGene/Q “Turing”. The graph shows the runtime needed for 10 ms of propagated activity using an explicit (forward) Euler integration. Dashed lines indicate the ideal scaling line with respect to the lowest number of cores measured. **Right:** Bundle structures and different layers of fiber orientation created by the Gepetto software.

6.3.3. MUSIC

MUSIC is a multimodal platform for cardiac imaging developed by the imaging team at IHU LIRYC in collaboration with the Inria team Asclepios (<https://bil.inria.fr/fr/software/view/1885/tab>). It is based on the medInria software also developed by the Asclepios team. MUSIC is a cross-platform software for segmentation of medical imaging data, meshing, and ultimately also visualization of functional imaging data and model results.

Several members of the Carmen team use MUSIC for their work. The team also contributes a series of plugins for MUSIC through the IDAM project (section [6.2](#)).

CQFD Project-Team

6. New Software and Platforms

6.1. Package divclust

FUNCTIONAL DESCRIPTION

DIVCLUS-T is a divisive hierarchical clustering algorithm based on a monothetic bipartitional approach allowing the dendrogram of the hierarchy to be read as a decision tree. It is designed for numerical, categorical (ordered or not) or mixed data. Like the Ward agglomerative hierarchical clustering algorithm and the k-means partitioning algorithm, it is based on the minimization of the inertia criterion. However, it provides a simple and natural monothetic interpretation of the clusters. Indeed, each cluster is described by set of binary questions. The inertia criterion is calculated on all the principal components of PCAmix (and then on standardized data in the numerical case).

- Participants: Marie Chavent, Marc Fuentes
- Contact: Marie Chavent
- URL: <https://github.com/chavent/divclust>

6.2. Package ClustGeo

FUNCTIONAL DESCRIPTION

This R package is dedicated to the clustering of objects with geographical positions. The clustering method implemented in this package allows the geographical constraints of proximity to be taken into account within the ascendant hierarchical clustering.

- Marie Chavent, Amaury Labenne, Vanessa Kuentz, Jérôme Saracco
- Contact: Amaury Labenne
- URL: <https://cran.r-project.org/web/packages/ClustGeo/index.html>

6.3. PCAmixdata

FUNCTIONAL DESCRIPTION

Mixed data type arise when observations are described by a mixture of numerical and categorical variables. The R package PCAmixdata extends standard multivariate analysis methods to incorporate this type of data. The key techniques included in the package are PCAmix (PCA of a mixture of numerical and categorical variables), PCArot (rotation in PCAmix) and MFAmix (multiple factor analysis with mixed data within a dataset). The MFAmix procedure handles a mixture of numerical and categorical variables within a group - something which was not possible in the standard MFA procedure. We also included techniques to project new observations onto the principal components of the three methods in the new version of the package.

- Participants: Marie Chavent, Amaury Labenne, Jérôme Saracco
- Contact: Marie Chavent
- URL: <https://cran.r-project.org/web/packages/PCAmixdata/index.html>

6.4. QuantifQuantile

FUNCTIONAL DESCRIPTION

QuantifQuantile is an R package that allows to perform quantization-based quantile regression. The different functions of the package allow the user to construct an optimal grid of N quantizers and to estimate conditional quantiles. This estimation requires a data driven selection of the size N of the grid that is implemented in the functions. Illustration of the selection of N is available, and graphical output of the resulting estimated curves or surfaces (depending on the dimension of the covariate) is directly provided via the plot function.

- Isabelle Charlier, Jérôme Saracco
- Contact: Isabelle Charlier
- URL: <https://cran.r-project.org/web/packages/QuantifQuantile/index.html>

6.5. biips

Bayesian Inference with Interacting Particle Systems

FUNCTIONAL DESCRIPTION

Biips is a software platform for automatic Bayesian inference with interacting particle systems. Biips allows users to define their statistical model in the probabilistic programming BUGS language, as well as to add custom functions or samplers within this language. Then it runs sequential Monte Carlo based algorithms (particle filters, particle independent Metropolis-Hastings, particle marginal Metropolis-Hastings) in a black-box manner so that to approximate the posterior distribution of interest as well as the marginal likelihood. The software is developed in C++ with interfaces with the softwares R, Matlab and Octave.

- Participants: Francois Caron and Adrien Todeschini
- Contact: Adrien Todeschini
- URL: <http://biips.gforge.inria.fr>

6.6. VCN: Software for analysis of VCN

FUNCTIONAL DESCRIPTION

VCN is a software for the analysis of the vigilance of the patient based on the analysis of the EEG signals. The code is written in Matlab and provides an interface easy to use for someone without informatics skills.

- Participants: Pierrick Legrand, Julien Clauzel, Laurent Vezard, Charlotte Rodriguez, Borjan Geshkovski.
- Contact: Pierrick Legrand

6.7. EMGView: Software for visualisation and time-frequency analysis of bio signals

FUNCTIONAL DESCRIPTION

EMGView is a software for the visualisation and the analysis of bio-signals. The code is written in Matlab and provides an interface easy to use for someone without informatics skills.

- Participants: Luis Herrera, Eric Grivel, Pierrick Legrand, Gregory Barriere
- Contact: Pierrick Legrand

FLOWERS Project-Team

6. New Software and Platforms

6.1. Poppy project

6.1.1. HiPi Board

FUNCTIONAL DESCRIPTION

HiPi is a board to control robots on Raspberry Pi. It is an extension of the Pixl board with the following features:

- A DC/DC power converter from 12V (motor) to 5V (Raspberry Pi) at 3A.
- A stereo audio amplifier 3W.
- A MPU9250 central motion unit .
- A RS232 and a RS485 bus connected to the Raspberry Pi by SPI for driving MX and RX Dynamixel motor series.

This board will be integrated soon in the new head of the Poppy Humanoid and Poppy Torso.

Using the Raspberry Pi for every Poppy robots will simplify the hardware complexity (we maintain 4 types of embedded boards, with different Linux kernel and configurations) and improve the usage and installation of new robots.

- Contact: Theo Segonds
- URL: <https://forum.poppy-project.org/t/poppy-1-1-hipi/2137>

6.1.2. IKPy

Inverse Kinematics Python Library

FUNCTIONAL DESCRIPTION

IKPy is a Python Inverse Kinematics library, designed to be simple to use and extend. It provides Forward and Inverse kinematics functionality, bundled with helper tools such as 3D plotting of the kinematics chains. Being written entirely in Python, IKPy is lightweight and is based on numpy and scipy for fast optimization. IKPy is compatible with many robots, by automatically parsing URDF files. It also supports other (such as DH-parameters) and custom representations. Moreover, it provides a framework to easily implement new Inverse Kinematics strategies. Originally developed for the Poppy project, it can also be used as a standalone library.

- Contact: Pierre Manceron
- URL: <https://github.com/Phylliade/ikpy>

6.1.3. Pixl Board

FUNCTIONAL DESCRIPTION

Pixl is a tiny board used to create low cost robots based on Raspberry Pi board and Dynamixel XL-320 motors. This board has 2 main features:

- The power part, allowing the user to plug a 7.5V AC/DC converter or a battery directly into the Pixl. This power is distributed to all XL320 motors and is converted to 5V for the Raspberry Pi board.
- The communication part, which converts full duplex to half duplex and vice-versa. The half duplex part switch between RX and TX automatically. Another connector allows the user to connect his XL320 network.

The board is used in the Poppy Ergo Jr robot.

- Contact: Theo Segonds
- URL: <https://github.com/poppy-project/pixl>

6.1.4. Poppy

FUNCTIONAL DESCRIPTION

The Poppy Project team develops open-source 3D printed robots platforms based on robust, flexible, easy-to-use and reproduce hardware and software. In particular, the use of 3D printing and rapid prototyping technologies is a central aspect of this project, and makes it easy and fast not only to reproduce the platform, but also to explore morphological variants. Poppy targets three domains of use: science, education and art.

In the Poppy project we are working on the Poppy System which is a new modular and open-source robotic architecture. It is designed to help people create and build custom robots. It permits, in a similar approach as Lego, building robots or smart objects using standardized elements.

Poppy System is a unified system in which essential robotic components (actuators, sensors...) are independent modules connected with other modules through standardized interfaces:

- Unified mechanical interfaces, simplifying the assembly process and the design of 3D printable parts.
- Unified communication between elements using the same connector and bus for each module.
- Unified software, making it easy to program each module independently.

Our ambition is to create an ecosystem around this system so communities can develop custom modules, following the Poppy System standards, which can be compatible with all other Poppy robots.

- Participants: Pierre Rouanet, Matthieu Lapeyre, Jonathan Grizou and Pierre-Yves Oudeyer
- Contact: Pierre-Yves Oudeyer
- URL: <https://www.poppy-project.org/>

6.1.5. Poppy Ergo Jr

FUNCTIONAL DESCRIPTION

Poppy Ergo Jr is an open hardware robot developed by the Poppy Project to explore the use of robots in classrooms for learning robotic and computer science.

It is available as a 6 or 4 degrees of freedom arm designed to be both expressive and low-cost. This is achieved by the use of FDM 3D printing and low cost Robotis XL-320 actuators. A Raspberry Pi camera is attached to the robot so it can detect object, faces or QR codes.

The Ergo Jr is controlled by the Pypot library and runs on a Raspberry pi 2 or 3 board. Communication between the Raspberry Pi and the actuators is made possible by the Pixl board we have designed.

The Poppy Ergo Jr robot has several 3D printed tools extending its capabilities. There are currently the lampshade, the gripper and a pen holder.

With the release of a new Raspberry Pi board early 2016, the Poppy Ergo Jr disk image was updated to support Raspberry Pi 2 and 3 boards. The disk image can be used seamlessly with a board or the other.

- Contact: Theo Segonds
- URL: <https://github.com/poppy-project/poppy-ergo-jr>

6.1.6. Poppy Ergo Jr Installer

FUNCTIONAL DESCRIPTION

An alternative way to install the Ergo Jr robot software is made available using containers.

Users can own their own operating system installation, then add the Ergo Jr required software in a sandboxed environment. This results in a non-intrusive installation on the host system.




Figure 1. Poppy Ergo Jr, 6-DoFs arm robot for education




Figure 2. The available Ergo Jr tools: a pen holder, a lampshade and a gripper

Docker containers implementation were used, and image is hosted at Docker Hub.

- Contact: Damien Caselli
- URL: <https://hub.docker.com/r/poppycommunity/ergo-jr/>

6.1.7. Poppy Ergo Jr Simulator

FUNCTIONAL DESCRIPTION

Poppy Project, through Poppy Education, wants users to get used to robotics, even without owning a physical robot.

For that purpose, Poppy Project team created a dummy robot in Pypot that is meant to be used in conjunction with a consumer application. We choose to develop a web hosted application using a 3D engine (Threejs) to render the robot.

Our ambition is to have a completely standalone simulated robot with physics. Some prototypes were created to benchmark possible solutions.

- Contact: Damien Caselli
- URL: <https://github.com/poppy-project/poppy-simu>

6.1.8. PyPot

SCIENTIFIC DESCRIPTION

Pypot is a framework developed to make it easy and fast to control custom robots based on Dynamixel motors. This framework provides different levels of abstraction corresponding to different types of use. Pypot can be used to:

- control Robotis motors through a USB2serial device,
- define the structure of a custom robot and control it through high-level commands,
- define primitives and easily combine them to create complex behavior.

Pypot is part of the Poppy project. It is the core library used by the Poppy robots. This abstraction layer allows to seamlessly switch from a given Poppy robot to another. It also provides a common set of tools, such as forward and inverse kinematics, simple computer vision, recording and replaying moves, or easy access to the autonomous exploration library Explauto.

To extend pypot application domains and connection to outside world, it also provides an HTTP API. On top of providing an easy way to connect to smart sensors or connected devices, it is notably used to connect to Snap!, a variant of the well-known Scratch visual programming language.

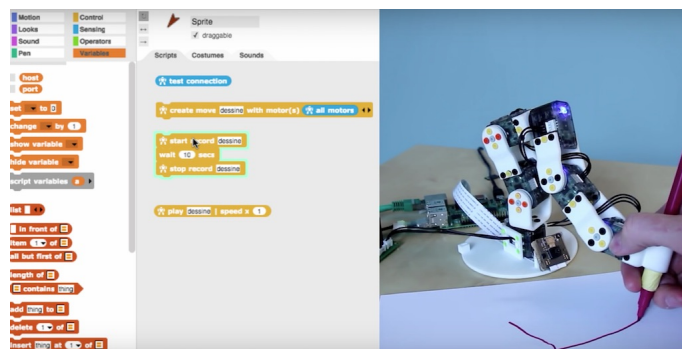


Figure 3. Example of using pypot to program a robot to reproduce a drawn shape

FUNCTIONAL DESCRIPTION

Pypot is entirely written in Python to allow for fast development, easy deployment and quick scripting by non-expert developers. It can also benefit from the scientific and machine learning libraries existing in Python. The serial communication is handled through the standard library and offers high performance (10ms sensorimotor loop) for common Poppy uses. It is cross-platform and has been tested on Linux, Windows and Mac OS.

Pypot is also compatible with the V-REP simulator. This allows the transparent switch from a real robot to its simulated equivalent with a single code base.

Finally, it has been developed to be easily and quickly extended for other types of motors and sensors.

It works with Python 2.7 or Python 3.3 or later, and has also been adapted to the Raspberry Pi board.

Pypot has been connected to Snap!, a variant of the famous Scratch visual language, developed to teach computer science to children. It is based on a drag-and-drop blocks interface to write scripts by assembling those blocks.

Thanks to the Snap! HTTP block, a connection can be made to pypot allowing users to directly control robots through their visual interfaces. A set of dedicated Snap! blocks have been designed, such as *set motor position* or *get motor temperature*. Thanks to the Snap! HTTP block, users can control robots through this visual interfaces connecting to Pypot. A set of dedicated Snap! blocks has been designed, such as *set motor position* or *get motor temperature*.



Figure 4. Using Snap! to program a robot by demonstration and create complex choreographies

Snap! is also used as a tool to program the robot by demonstration. Using the *record* and *play* blocks, users can easily trigger kinesthetic recording of the whole robot or only a specific subpart, such as an arm. These records can then be played or "mixed" - either played in sequence or simultaneously - with other recordings to compose complex choreographies. The moves are encoded as a model of mixture of gaussians (GMM) which allows the definition of clean mathematical operators for combining them.

This recording tool has been developed and used in collaboration with artists who show interest in the concept of robotic moves.

- Participants: Pierre Rouanet, Matthieu Lapeyre, Steve Nguyen, Damien Caselli and Theo Segonds
- Contact: Theo Segonds
- URL: <https://github.com/poppy-project/pypot>




Figure 5. Artistic project exploring the concept of robotic move.

6.1.9. PyQMC

Python library for Quasi-Metric Control
FUNCTIONAL DESCRIPTION

PyQMC is a python library implementing the control method described in <http://dx.doi.org/10.1371/journal.pone.0083411>. It allows to solve discrete markovian decision processes by computing a Quasi-Metric on the state space. This model based method has the advantage to be goal independant and thus can produce a policy for any goal with relatively few recomputation. New addition to this method is the possibility of online learning of the transition model and the Quasi-Metric.

- Participant: Steve Nguyen
- Contact: Steve Nguyen
- URL: <https://github.com/SteveNguyen/pyqmc>

6.2. Explauto

an autonomous exploration library
SCIENTIFIC DESCRIPTION

An important challenge in developmental robotics is how robots can be intrinsically motivated to learn efficiently parametrized policies to solve parametrized multi-task reinforcement learning problems, i.e. learn the mappings between the actions and the problem they solve, or sensory effects they produce. This can be a robot learning how arm movements make physical objects move, or how movements of a virtual vocal tract modulates vocalization sounds. The way the robot will collect its own sensorimotor experience have a strong impact on learning efficiency because for most robotic systems the involved spaces are high dimensional, the mapping between them is non-linear and redundant, and there is limited time allowed for learning. If robots explore the world in an unorganized manner, e.g. randomly, learning algorithms will be often ineffective because very sparse data points will be collected. Data are precious due to the high dimensionality and the limited time, whereas data are not equally useful due to non-linearity and redundancy. This is why learning has to be guided using efficient exploration strategies, allowing the robot to actively drive its own interaction with the environment in order to gather maximally informative data to optimize the parametrized policies. In the recent year, work in developmental learning has explored various families of algorithmic principles which allow the efficient guiding of learning and exploration.

Explauto is a framework developed to study, model and simulate curiosity-driven learning and exploration in real and simulated robotic agents. Explauto's scientific roots trace back from Intelligent Adaptive Curiosity algorithmic architecture [152], which has been extended to a more general family of autonomous exploration architectures by [3] and recently expressed as a compact and unified formalism [38]. The library is detailed in [39]. In Explauto, interest models are implementing the strategies of active selection of particular problems / goals in a parametrized multi-task reinforcement learning setup to efficiently learn parametrized policies. The agent can have different available strategies, parametrized problems, models, sources of information, or learning mechanisms (for instance imitate by mimicking vs by emulation, or asking help to one teacher or to another), and chooses between them in order to optimize learning (a process called strategic learning [45]). Given a set of parametrized problems, a particular exploration strategy is to randomly draw goals/ RL problems to solve in the motor or problem space. More efficient strategies are based on the active choice of learning experiments that maximize learning progress using bandit algorithms, e.g. maximizing improvement of predictions or of competences to solve RL problems [152]. This automatically drives the system to explore and learn first easy skills, and then explore skills of progressively increasing complexity. Both random and learning progress strategies can act either on the motor or on the problem space, resulting in motor babbling or goal babbling strategies.

- Motor babbling consists in sampling commands in the motor space according to a given strategy (random or learning progress), predicting the expected effect, executing the command through the environment and observing the actual effect. Both the parametrized policies and interest models are finally updated according to this experience.
- Goal babbling consists in sampling goals in the problem space and to use the current policies to infer a motor action supposed to solve the problem (inverse prediction). The robot/agent then executes the command through the environment and observes the actual effect. Both the parametrized policies and interest models are finally updated according to this experience. It has been shown that this second strategy allows a progressive solving of problems much more uniformly in the problem space than with a motor babbling strategy, where the agent samples directly in the motor space [3].




Figure 6. Complex parametrized policies involve high dimensional action and effect spaces. For the sake of visualization, the motor M and sensory S spaces are only 2D each in this example. The relationship between M and S is non-linear, dividing the sensorimotor space into regions of unequal stability: small regions of S can be reached very precisely by large regions of M , or large regions in S can be very sensitive to variations in M : s as well as a non-linear and redundant relationship. This non-linearity can imply redundancy, where the same sensory effect can be attained using distinct regions in M .

FUNCTIONAL DESCRIPTION

This library provides high-level API for an easy definition of:

- Real and simulated robotic setups (Environment level),
- Incremental learning of parametrized policies (Sensorimotor level),
- Active selection of parametrized RL problems (Interest level).

The library comes with several built-in environments. Two of them corresponds to simulated environments: a multi-DoF arm acting on a 2D plan, and an under-actuated torque-controlled pendulum. The third one allows to control real robots based on Dynamixel actuators using the Pypot library. Learning parametrized policies involves machine learning algorithms, which are typically regression algorithms to learn forward models, from motor controllers to sensory effects, and optimization algorithms to learn inverse models, from sensory effects, or problems, to the motor programs allowing to reach them. We call these sensorimotor learning algorithms sensorimotor models. The library comes with several built-in sensorimotor models: simple nearest-neighbor look-up, non-parametric models combining classical regressions and optimization algorithms, online mixtures of Gaussians, and discrete Lidstone distributions. Explauto sensorimotor models are online learning algorithms, i.e. they are trained iteratively during the interaction of the robot in the environment in which it evolves. Explauto provides also a unified interface to define exploration strategies using the InterestModel class. The library comes with two built-in interest models: random sampling as well as sampling maximizing the learning progress in forward or inverse predictions.

Explauto environments now handle actions depending on a current context, as for instance in an environment where a robotic arm is trying to catch a ball: the arm trajectories will depend on the current position of the ball (context). Also, if the dynamic of the environment is changing over time, a new sensorimotor model (Non-Stationary Nearest Neighbor) is able to cope with those changes by taking more into account recent experiences. Those new features are explained in Jupyter notebooks.

This library has been used in many experiments including:

- the control of a 2D simulated arm,
- the exploration of the inverse kinematics of a poppy humanoid (both on the real robot and on the simulated version),
- acoustic model of a vocal tract.

Explauto is cross-platform and has been tested on Linux, Windows and Mac OS. It has been released under the GPLv3 license.

- Contact: Sébastien Forestier
- URL: <https://github.com/flowersteam/explauto>

6.3. Tools for robot learning, control and perception

6.3.1. CARROMAN

FUNCTIONAL DESCRIPTION

This software implements a control architecture for the Meka humanoid robot. It integrates the Stanford Whole Body Control in the M3 architecture provided with the Meka robot, and provides clear and easy to use interfaces through the URBI scripting language. This software provides a modular library of control modes and basic skills for manipulating objects, detecting objects and humans which other research projects can reuse, extend and enhance. An example would be to locate a cylindrical object on a table using stereo vision, and grasping it using position and force control.

- Contact: David Filliat

6.3.2. Aversive++

FUNCTIONAL DESCRIPTION

Aversive++ is a C++ library that eases micro-controller programming. Its aim is to provide an interface simple enough to be able to create complex applications, and optimized enough to enable small micro-controllers to execute these applications. The other aspect of this library is to be multiplatform. Indeed, it is designed to provide the same API for a simulator (named SASIAE) and for AVR-based and ARM-based micro-controllers.

- Contact: Loïc Dauphin
- URL: <http://aversiveplusplus.com/>

6.3.3. DMP-BBO

Black-Box Optimization for Dynamic Movement Primitives

KEYWORD: -

FUNCTIONAL DESCRIPTION

The DMP-BBO Matlab library is a direct consequence of the insight that black-box optimization outperforms reinforcement learning when using policies represented as Dynamic Movement Primitives. It implements several variants of the PIBB algorithm for direct policy search. The dmp-bbo C++ library has been extended to include the “unified model for regression”. The implementation of several of the function approximators have been made real-time compatible.

- Participant: Freek Stulp
- Partner: ENSTA
- Contact: Freek Stulp
- URL: <https://github.com/stulp/dmpbbo>

6.3.4. KERAS-QR

KERAS with Quick Reset

KEYWORDS: Library - Deep learning

- Participant: Florian Golemo
- Contact: Florian Golemo
- URL: <https://github.com/fgolemo/keras>

6.3.5. Multimodal

FUNCTIONAL DESCRIPTION

The python code provides a minimum set of tools and associated libraries to reproduce the experiments in [98], together with the choreography datasets. The code is primarily intended for reproduction of the multimodal learning experiment mentioned above. It has already been reused in several experimentations by other member of the team and is expected to play an important role in further collaborations. It is also expected that the public availability of the code encourages further experimentation by other scientists with data coming from other domains, thus increasing both the impact of the aforementioned publication and the knowledge on the algorithm behaviors.

- Participant: Olivier Mangin
- Contact: Olivier Mangin
- URL: <https://github.com/omangin/multimodal>

6.3.6. Of 3-D point cloud

FUNCTIONAL DESCRIPTION

This software scans the 3-D point cloud of a scene to find objects and match them against a database of known objects. The process consists in 3 stages. The segmentation step finds the objects in the point cloud, the feature extraction computes discriminating properties to be used in the classification stage for object recognition.

- Participants: David Filliat, Alexander Geppert and Louis-Charles Caron
- Contact: Alexander Geppert

6.3.7. PEDDETECT

FUNCTIONAL DESCRIPTION

PEDDETECT implements real-time person detection in indoor or outdoor environments. It can grab image data directly from one or several USB cameras, as well as from pre-recorded video streams. It detects multiple persons in 800x600 color images at frame rates of >15Hz, depending on available GPU power. In addition, it also classifies the pose of detected persons in one of the four categories "seen from the front", "seen from the back", "facing left" and "facing right". The software makes use of advanced feature computation and nonlinear SVM techniques which are accelerated using the CUDA interface to GPU programming to achieve high frame rates. It was developed in the context of an ongoing collaboration with Honda Research Institute USA, Inc.

- Participant: Alexander Gepperth
- Contact: Alexander Gepperth

6.3.8. ThifloNet

KEYWORDS: Deep learning - Policy Learning

SCIENTIFIC DESCRIPTION

We created a software architecture that combines a state-of-the-art computer vision system with a policy learning framework. This system is able to perceive a visual scene, given by a still image, extract facts ("predicates"), and propose an optimal action to achieve a given goal. Both systems are chained into a pipeline that is trained by presenting images and demonstrating an optimal action. By providing this information, both the predicate recognition model and the policy learning model are updated.

Our architecture is based on the recent works of Lerer, A., Gross, S., & Fergus, R., 2016 ("Learning Physical Intuition of Block Towers by Example"). They created a large network able to identify physical properties of stacked blocks. Analogously our vision system utilizes the same network layout (without the image prediction auxiliary output), with an added output layer for predicates, based on the expected number and arity of predicates. The vision subsystem is not trained with a common cross-entropy or MSE loss function, but instead receives its loss from the policy learning subsystem. The policy learning module calculates the loss as optimal combination of predicates for the given expert action.

By using this combination of systems, the architecture as a whole requires significantly fewer data samples than other systems (which exclusively utilize neural networks). This makes the approach more feasible to real-life application with actual live demonstration.

FUNCTIONAL DESCRIPTION

The neural network consists of ResNet-50 (the currently best-performing computer vision system), with 50 layers, 2 layers for converting the output of ResNet to predicates and a varying amount of output neurons, corresponding to the estimated number of n-arity predicates. The network was pretrained on the ImageNet dataset. The policy learning module incorporates the ACE tree learning tool and a wrapper in Prolog.

Our example domain consists of 2-4 cubes colored in red, blue, green, and yellow and randomly stacked on top of each other in a virtual 3D environment. The dataset used for training and testing contains a total of 30000 elements, each with an image of the scene, the correct predicates, a list of blocks that are present and the corresponding expert action, that would lead to stacking the blocks to a tower.

- Participants: Florian Golemo, Thibaut Munzer and Manuel Lopes
- Contact: Florian Golemo

6.4. Tools for education

6.4.1. KidLearn

KEYWORD: Automatic Learning

FUNCTIONAL DESCRIPTION

KidLearn is a software which adaptively personalizes sequences of learning activities to the particularities of each individual student. It aims at proposing to the student the right activity at the right time, maximizing concurrently his learning progress and its motivation. The library regrouping the different developed technologies is available on [github](#).

- Participants: Benjamin Clement, Pierre Yves Oudeyer, Didier Roy and Manuel Lopes
- Contact: Manuel Lopes
- URL: <https://flowers.inria.fr/research/kidlearn/>, <https://github.com/flowersteam/kidlearn>

6.4.2. KidBreath

FUNCTIONAL DESCRIPTION

KidBreath is a web responsive application composed by several interactive contents linked to asthma and displayed to different forms: learning activities with quiz, short games and videos. There are profil creation and personalization, and a part which describes historic and scoring of learning activities, to see evolution of Kidreath use. To test Kidlearn algorithm, it is adapted and integrated on this platform. Development in PHP, HTML-5, CSS, MySQL, JQuery, Javascript. Hosting in APACHE, LINUX, PHP 5.5, MySQL, OVH.

- Partner: ItWell SAS
- Contact: Alexandra Delmas
- URL: <http://www.kidbreath.fr>

6.4.3. Kidlearn: money game application

FUNCTIONAL DESCRIPTION

The game is instantiated in a browser environment where students are proposed exercises in the form of money/token games (see Figure 7). For an exercise type, one object is presented with a given tagged price and the learner has to choose which combination of bank notes, coins or abstract tokens need to be taken from the wallet to buy the object, with various constraints depending on exercises parameters. The games have been developed using web technologies, HTML5, javascript and Django.

- Contact: Benjamin Clement
- URL: <https://flowers.inria.fr/research/kidlearn/>

6.4.4. Kidlearn: script for Kidbreath use

FUNCTIONAL DESCRIPTION

A new way to test Kidlearn algorithms is to use them on Kidbreath Platform. The Kidbreath Platform use apache/PHP server, so to facilitate the integration of our algorithm, a python script have been made to allow PHP code to use easily the python library already made which include our algorithms.

Github link to explanation about it : https://github.com/flowersteam/kidlearn/tree/feature/kidbreath/module_php.

- Contact: Benjamin Clement
- URL: https://github.com/flowersteam/kidlearn/tree/feature/kidbreath/module_php

6.5. 3rdHand Project

6.5.1. 3rdHand Infrastructure

KEYWORDS: Interaction - Robotics - Infrastructure software - Framework - Robot Operating System (ROS)

FUNCTIONAL DESCRIPTION

The infrastructure is predicate-based to handle relational actions and covers perception (scene description generation, human actions recognition), decision making (teleoperated, scripted or learning from demonstrations), interaction with end users (GUI, voice, gestures) and parallel executions of robotic actions (hold, pick, grasp, bring, ...).

- Contact: Yoan Mollard
- https://github.com/3rdHand-project/thr_infrastructure

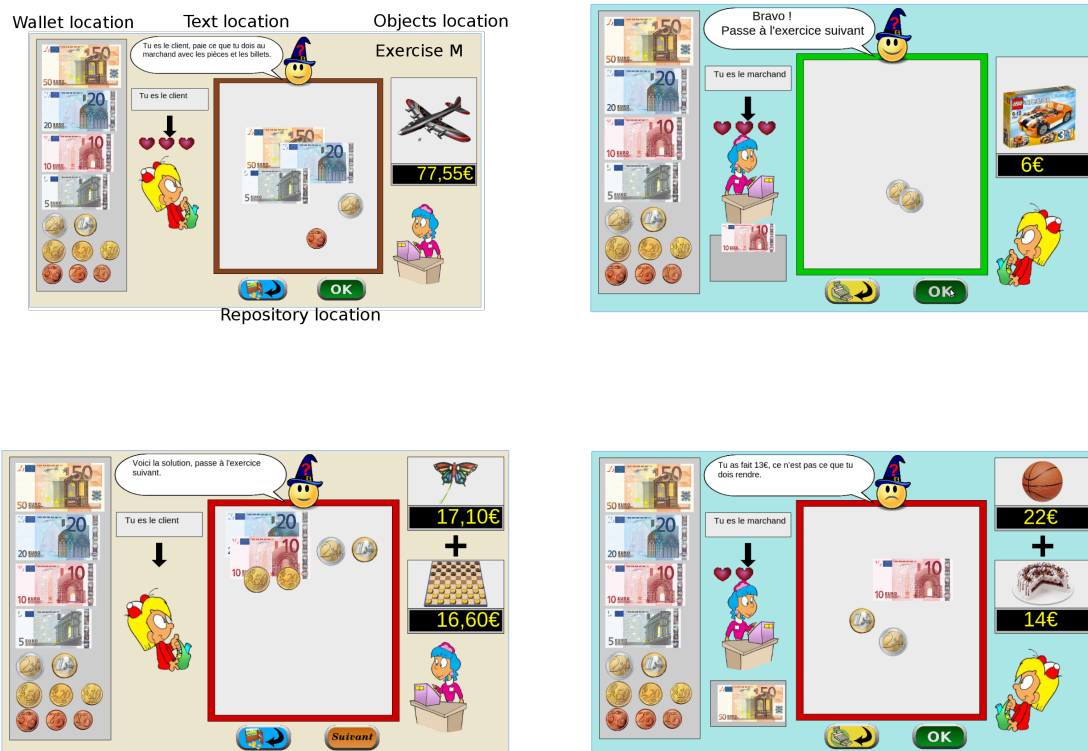


Figure 7. Four principal regions are defined in the graphical interface. The first is the wallet location where users can pick and drag the money items and drop them on the repository location to compose the correct price. The object and the price are present in the object location. Four different types of exercises exist: M : customer/one object, R : merchant/one object, MM : customer/two objects, RM : merchant/two objects.

6.5.2. Kinect 2 Server

Kinect 2 server

KEYWORDS: Depth Perception - Speech recognition - Gesture recognition - Kinect

FUNCTIONAL DESCRIPTION

The server written in C# uses the Kinect SDK v2 to get the RGBD raw image, skeleton tracking information, recognized speech. It also uses the text-to-speech from Microsoft. Then it streams JSON data over the network using the Publisher/Subscriber pattern from the ZeroMQ network library. A Linux client has been written in Python but it can be written in any other language that is compatible with ZeroMQ. Features are controllable through a Graphical User Interface on Windows, or through the code from any Linux/Windows client. The clients can for instance enable features (speech recognition on, skeleton tracking off, ...) and parameters (set new speech to recognize, change language, ...) from remote.

- Contact: Yoan Mollard
- URL: https://github.com/baxter-flowers/kinect_2_server/

6.5.3. ProMP

Probabilistic Movement Primitives

KEYWORDS: Interaction - Robotics - Probability - Motion model - Robot Operating System (ROS)

FUNCTIONAL DESCRIPTION

Joint-space primitives with a task-space constraint: The primitives are stored in joint-space but demonstrations are provided both in joint space and task space, context. Thanks to this context, task-space goals can be requested to these joint-space primitives. The benefit is that requesting a new task-space goal does not require to call an IK method which would return demonstrations-agnostic joint configurations.

Vocal interactive learning and clustering: This work includes an interactive learning aspect which allows to automatically cluster motor primitives based on the standard deviation of their demonstrations. A new primitive is created automatically if the provided demonstration is out of 2 standard deviation of the existing primitives, otherwise the demonstration is distributed to an existing one.

- Contact: Yoan Mollard
- URL: <https://github.com/baxter-flowers/promplib>

6.5.4. ROS Optitrack Publisher

KEYWORDS: Target tracking - Robot Operating System (ROS)

FUNCTIONAL DESCRIPTION

This package allows to publish optitrack markers declared as rigid bodies as TF transforms. Data is gathered through the embedded VRPN server of Motive/Arena. Only rigid bodies are requested to the server, thus single points in 2D/3D are ignored. VRPN server can be enable in View > Data streaming in Motive.

- Contact: Yoan Mollard
- URL: https://github.com/baxter-flowers/optitrack_publisher

GEOSTAT Project-Team

5. New Software and Platforms

5.1. Fluex

KEYWORDS: Signal - Signal processing

SCIENTIFIC DESCRIPTION

Fluex is a package consisting of the Microcanonical Multiscale Formalism for 1D, 2D 3D and 3D+t general signals.

FUNCTIONAL DESCRIPTION Fluex is a C++ library developed under Gforge. Fluex is a library in nonlinear signal processing. Fluex is able to analyze turbulent and natural complex signals, Fluex is able to determine low level features in these signals that cannot be determined using standard linear techniques.

- Participants: Rémi Paties and Hussein Yahia
- Contact: Hussein Yahia
- URL: <https://geostat.bordeaux.inria.fr/index.php/downloads.html>
- URL: <https://bil.inria.fr/fr/software/view/2113/tab>

5.2. FluidExponents

- Participants: Hussein Yahia and Antonio Turiel
- Contact: Hussein Yahia
- URL: <https://bil.inria.fr/fr/software/view/336/tab>

HIEPACS Project-Team

6. New Software and Platforms

6.1. Spack-morse

The radical change we have adopted in terms of methodology (task-based programming strongly) changes the software design. In particular, our codes become more and more modular and the complexity of their inter-dependencies is subsequently very high.

In order to address this complexity we have chosen to rely on the **Spack** flexible package manager designed to support multiple versions, configurations, platforms, and compilers (<http://software.llnl.gov/spack> developed and maintained at LLNL). We have integrated all our libraries above this package manager in the **Spack-Morse** extension that we maintain in **HIEPACS**.

- Audience: A-4 (large audience, used by people outside the team).
- Software originality: SO-3 (original software reusing known ideas and introducing new ideas).
- Software maturity: SM-3 (well-developed software, good documentation, reasonable software engineering).
- Evolution and maintenance: EM-3 (good quality middle-term maintenance).
- Software distribution and licensing: SDL-4 (public source or binary distribution on the Web). source distribution or a commercially-distributed product).
- Contact: Florent Pruvost
- URL: <http://morse.gforge.inria.fr/spack/spack.html>

6.2. Chameleon

Chameleon is a dense linear algebra software relying on the STF sequential task-based programming paradigm. It implements the tile algorithms originally designed for multicore architectures in the **PLASMA** package and extends them so that they can be processed on by a runtime system to exploit any type of hardware architecture (multicore, GPU, heterogeneous, supercomputer). This software is central for the team as it allows to investigate in a relatively simple context (regular dense linear algebra algorithms) new types of designs before implementing them for the more irregular algorithms implemented in the software packages described below.

- Audience: A-4 (large audience, used by people outside the team).
- Software originality: SO-4 (original software implementing a fair number of original ideas).
- Software maturity: SM-3 (well-developed software, good documentation, reasonable software engineering).
- Evolution and maintenance: EM-3 (good quality middle-term maintenance).
- Software distribution and licensing: SDL-4 (public source or binary distribution on the Web). source distribution or a commercially-distributed product).
- Contact: Emmanuel Agullo
- URL: <https://project.inria.fr/chameleon>

6.3. HIPS

HIPS (Hierarchical Iterative Parallel Solver) is a scientific library that provides an efficient parallel iterative solver for very large sparse linear systems.

The key point of the methods implemented in **HIPS** is to define an ordering and a partition of the unknowns that relies on a form of nested dissection ordering in which cross points in the separators play a special role (Hierarchical Interface Decomposition ordering). The subgraphs obtained by nested dissection correspond to the unknowns that are eliminated using a direct method and the Schur complement system on the remaining of the unknowns (that correspond to the interface between the sub-graphs viewed as sub-domains) is solved using an iterative method (GMRES or Conjugate Gradient at the time being).

Thus, **HIPS** is a software library that provides several methods to build an efficient preconditioner in almost all situations.

- Audience: A-4 (large audience, used by people outside the team).
- Software originality: SO-4 (original software implementing a fair number of original ideas).
- Software maturity: SM-3 (well-developed software, good documentation, reasonable software engineering).
- Evolution and maintenance: EM-2 (basic maintenance to keep the software alive).
- Software distribution and licensing: SDL-4 (public source or binary distribution on the Web).
- Contact: Pierre Ramet
- URL: <http://hips.gforge.inria.fr>

6.4. MaPHYS

MaPHyS (Massively Parallel Hybrid Solver) is an hybrid iterative/direct parallel (MPI-threads) sparse linear solver based on algebraic domain decomposition technique for real/complex symmetric positive definite/unsym/-metric matrices. For a given number of MPI processes/domains, **MaPHyS** solves the Schur complement computed from the adjacency graph of the sparse matrix using a preconditioned Krylov subspace method (CG or GMRES). The provided preconditioners are variants of an algebraic Additive Schwarz methods. A prototype version of a two level preconditionner using an algebraic coarse space is available but not yet publicly distributed (provided upon request for beta testers).

- Audience: A-4 (large audience, used by people outside the team).
- Software originality: SO-4 (original software implementing a fair number of original ideas).
- Software maturity: SM-3 (well-developed software, good documentation, reasonable software engineering).
- Evolution and maintenance: EM-4 (well-defined and implemented plans for future maintenance and evolution).
- Software distribution and licensing: SDL-4 (public source or binary distribution on the Web).
- Contact: Emmanuel Agullo
- URL: <https://project.inria.fr/maphys/fr>

6.5. MetaPart

MetaPart is a library that addresses the challenge of (dynamic) load balancing for emerging complex parallel simulations, such as multi-physics or multi-scale coupling applications. First, it offers a uniform API over state-of-the-art (hyper-) graph partitioning & ordering software packages such as **Scotch**, PaToH, METIS, Zoltan, Mondriaan, etc. Based upon this API, it provides a framework that facilitates the development and the evaluation of high-level partitioning methods, such as MxN repartitioning or coupling-aware partitioning (co-partitioning).

- Audience: A-1 (internal prototype).
- Software originality: SO-3 (original software reusing known ideas and introducing new ideas).
- Software maturity: SM-2 (basic usage works, terse documentation).
- Evolution and maintenance: EM-3 (good quality middle-term maintenance).
- Software distribution and licensing: SDL-4 (public source or binary distribution on the Web).
- Contact: Aurélien Esnard
- URL: <http://metapart.gforge.inria.fr>

6.6. PaStiX

PaStiX (Parallel Sparse matriX package) is a scientific library that provides a high performance parallel solver for very large sparse linear systems based on block direct and block ILU(k) iterative methods. Numerical algorithms are implemented in single or double precision (real or complex): LLt (Cholesky), LDLt (Crout) and LU with static pivoting (for non symmetric matrices having a symmetric pattern).

The **PaStiX** solver is suitable for any heterogeneous parallel/distributed architecture when its performance is predictable, such as clusters of multicore nodes. In particular, we now offer a high-performance version with a low memory overhead for multicore node architectures, which fully exploits the advantage of shared memory by using an hybrid MPI-thread implementation.

- Audience: A-5 (wide audience, large user's community).
- Software originality: SO-4 (original software implementing a fair number of original ideas).
- Software maturity: SM-4 (major software project, strong software engineering).
- Evolution and maintenance: EM-4 (well-defined and implemented plans for future maintenance and evolution).
- Software distribution and licensing: SDL-5 (external packaging and distribution, as part of a popular open source distribution or a commercially-distributed product).
- Contact: Pierre Ramet
- URL: <http://pastix.gforge.inria.fr>

6.7. QR_Mumps

qr_mumps is a software package for the solution of sparse, linear systems on multicore computers. It implements a direct solution method based on the QR factorization of the input matrix. Therefore, it is suited to solving sparse least-squares problems and to computing the minimum-norm solution of sparse, under-determined problems. It can obviously be used for solving square problems in which case the stability provided by the use of orthogonal transformations comes at the cost of a higher operation count with respect to solvers based on, e.g., the LU factorization. **qr_mumps** supports real and complex, single or double precision arithmetic.

qr_mumps is mainly developed and maintained by the APO team of the IRIT laboratory of Toulouse. **HIEPACS** is an active contributor to this project.

- Audience: A-4 (large audience, used by people outside the team).
- Software originality: SO-4 (original software implementing a fair number of original ideas).
- Software maturity: SM-3 (well-developed software, good documentation, reasonable software engineering).
- Evolution and maintenance: EM-3 (good quality middle-term maintenance).
- Software distribution and licensing: SDL-4 (public source or binary distribution on the Web).
- Contact: Emmanuel Agullo
- URL: http://buttari.perso.enseeiht.fr/qr_mumps/

6.8. ScalFMM

ScalFMM is a library to compute N-body interactions using the Fast Multipole Method. This is a parallel kernel independent fast multipole method based on interpolation (Chebychev or equispaced grid points).

ScalFMM intends to offer all the functionalities needed to perform large parallel simulations while enabling an easy customization of the simulation components: kernels, particles and cells. It works in parallel in a shared/distributed memory model using OpenMP (fork-join and tasks models), MPI and runtime system (**StarPU**). The software architecture has been designed with two major objectives: being easy to maintain and easy to understand. There is two main parts:

- the management of the tree structure (hierarchical octree and Group-Tree) and the parallel algorithms ;
- the kernels (scalar, tensorial and multi-rhs). Classical kernels are available (Coulombic, Leonard-Jones, Gaussian, Stokes, ...)

This modular architecture allows us to easily add new FMM algorithms or kernels and new paradigm of parallelization. Today, we also proposed the FMM based on spherical harmonic expansion with Blas or rotation optimization for Coulombic potential) and all algorithms are designed to treat more complex kernels by adding multiple right-hand sides, tensorial structures, ...

- Audience: A-4 (large audience, used by people outside the team).
- Software originality: SO-4 (original software implementing a fair number of original ideas).
- Software maturity: SM-3 (well-developed software, good documentation, reasonable software engineering).
- Evolution and maintenance: EM-3 (good quality middle-term maintenance).
- Software distribution and licensing: SDL-4 (public source or binary distribution on the Web).
- Contact: Olivier Coulaud
- URL: <http://scalfmm-public.gforge.inria.fr/doc/>

6.9. ViTE

ViTE is a trace explorer. It is a tool to visualize execution traces in Pajé or OTF format for debugging and profiling parallel or distributed applications. It is developed with C++ programming language with OpenGL and Qt technologies.

- Audience: A-4 (large audience, used by people outside the team).
- Software originality: SO-3 (original software reusing known ideas and introducing new ideas).
- Software maturity: SM-3 (well-developed software, good documentation, reasonable software engineering).
- Evolution and maintenance: EM-2 (basic maintenance to keep the software alive).
- Software distribution and licensing: SDL-4 (public source or binary distribution on the Web).
- Contact: Mathieu Faverge
- URL: <http://vite.gforge.inria.fr>

6.10. Platforms

6.10.1. PlaFRIM: Plateforme Fédérative pour la Recherche en Informatique et Mathématiques

PlaFRIM is an experimental platform for research in modeling, simulations and high performance computing. This platform has been set up from 2009 under the leadership of Inria Bordeaux Sud-Ouest in collaboration with computer science and mathematics laboratories, respectively Labri and IMB with a strong support in the region Aquitaine.

It aggregates different kinds of computational resources for research and development purposes. The latest technologies in terms of processors, memories and architecture are added when they are available on the market. It is now more than 1,000 cores (excluding GPU and Xeon Phi) that are available for all research teams of Inria Bordeaux, Labri and IMB. This computer is in particular used by all the engineers who work in HiePACS and are advised by F. Rue from the SED.

- Contact: Olivier Coulaud
- URL: <https://www.plafrim.fr/en/home/>

LFANT Project-Team

5. New Software and Platforms

5.1. APIP

Another Pairing Implementation in PARI

SCIENTIFIC DESCRIPTION

Apip , Another Pairing Implementation in PARI, is a library for computing standard and optimised variants of most cryptographic pairings.

The following pairings are available: Weil, Tate, ate and twisted ate, optimised versions (à la Vercauteren–Hess) of ate and twisted ate for selected curve families.

The following methods to compute the Miller part are implemented: standard Miller double-and-add method, standard Miller using a non-adjacent form, Boxall et al. version, Boxall et al. version using a non-adjacent form.

The final exponentiation part can be computed using one of the following variants: naive exponentiation, interleaved method, Avanzi–Mihalescu’s method, Kato et al.’s method, Scott et al.’s method.

Part of the library has been included into Pari/Gp proper.

FUNCTIONAL DESCRIPTION

APIP is a library for computing standard and optimised variants of most cryptographic pairings.

- Participant: Jérôme Milan
- Contact: Jérôme Milan
- URL: <http://www.lix.polytechnique.fr/~milanj/apip/apip.xhtml>

5.2. Arb

FUNCTIONAL DESCRIPTION

Arb is a C library for arbitrary-precision floating-point ball arithmetic. It supports real and complex numbers, polynomials, power series, matrices, and evaluation of many transcendental functions. All is done with automatic, rigorous error bounds. It has been accepted for inclusion in SageMath.

- Participant: Fredrik Johansson
- Contact: Fredrik Johansson
- URL: <http://fredrikj.net/arb/>

5.3. AVIsogenies

Abelian Varieties and Isogenies

FUNCTIONAL DESCRIPTION

AVIsogenies is a Magma package for working with abelian varieties, with a particular emphasis on explicit isogeny computation.

Its prominent feature is the computation of (l,l) -isogenies between Jacobian varieties of genus-two hyperelliptic curves over finite fields of characteristic coprime to l , practical runs have used values of l in the hundreds.

It can also be used to compute endomorphism rings of abelian surfaces, and find complete addition laws on them.

- Participants: Gaëtan Bisson, Romain Cosset and Damien Robert
- Contact: Damien Robert
- URL: <http://avisogenies.gforge.inria.fr/>

5.4. CM

FUNCTIONAL DESCRIPTION

The Cm software implements the construction of ring class fields of imaginary quadratic number fields and of elliptic curves with complex multiplication via floating point approximations. It consists of libraries that can be called from within a C program and of executable command line applications.

- Participant: Andreas Enge
- Contact: Andreas Enge
- URL: <http://www.multiprecision.org/index.php?prog=cm&page=home>

5.5. CMH

Computation of Igusa Class Polynomials

KEYWORDS: Mathematics - Cryptography - Number theory

FUNCTIONAL DESCRIPTION

Cmh computes Igusa class polynomials, parameterising two-dimensional abelian varieties (or, equivalently, Jacobians of hyperelliptic curves of genus 2) with given complex multiplication.

- Participants: Emmanuel Thomé, Andreas Enge and Regis Dupont
- Contact: Emmanuel Thomé
- URL: <http://cmh.gforge.inria.fr>

5.6. CUBIC

FUNCTIONAL DESCRIPTION

Cubic is a stand-alone program that prints out generating equations for cubic fields of either signature and bounded discriminant. It depends on the Pari library. The algorithm has quasi-linear time complexity in the size of the output.

- Participant: Karim Belabas
- Contact: Karim Belabas
- URL: <http://www.math.u-bordeaux.fr/~belabas/research/software/cubic-1.3.tgz>

5.7. Euclid

FUNCTIONAL DESCRIPTION

Euclid is a program to compute the Euclidean minimum of a number field. It is the practical implementation of the algorithm described in [38]. Some corresponding tables built with the algorithm are also available. Euclid is a stand-alone program depending on the PARI library.

- Participants: Pierre Lezowski and Jean-Paul Cerri
- Contact: Pierre Lezowski
- URL: <http://www.math.u-bordeaux1.fr/~plezowsk/euclid/index.php>

5.8. FLINT

FUNCTIONAL DESCRIPTION FLINT is a C library for number theory and basic computer algebra, maintained by William Hart with code by William Hart, Sebastian Pancratz, Andy Novocin, Fredrik Johansson, Tom Bachmann, Mike Hansen, Martin Lee, David Harvey, and a large number of other authors.

FLINT is used as a back end library for polynomial arithmetic and number theory functionality in a large number of applications, including SageMath and Singular.

- Participant: Fredrik Johansson
- Contact: William Hart
- URL: <http://flintlib.org/>

5.9. GNU MPC

FUNCTIONAL DESCRIPTION

Mpc is a C library for the arithmetic of complex numbers with arbitrarily high precision and correct rounding of the result. It is built upon and follows the same principles as Mpfr. The library is written by Andreas Enge, Philippe Théveny and Paul Zimmermann.

- Participants: Andreas Enge, Paul Zimmermann, Philippe Theveny and Mickaël Gastineau
- Contact: Andreas Enge
- URL: <http://www.multiprecision.org/>

5.10. KleinianGroups

FUNCTIONAL DESCRIPTION

KleinianGroups is a Magma package that computes fundamental domains of arithmetic Kleinian groups.

- Participant: Aurel Page
- Contact: Aurel Page
- URL: <http://www.normalesup.org/~page/Recherche/Logiciels/logiciels-en.html>

5.11. mpmath

FUNCTIONAL DESCRIPTION mpmath is a Python library for real and complex floating-point arithmetic with arbitrary precision. It has been developed by Fredrik Johansson since 2007, with help from many contributors.

As a dependency of the SymPy computer algebra system as well as SageMath, mpmath is a core component of the Python scientific software ecosystem.

- Participant: Fredrik Johansson
- Contact: Fredrik Johansson
- URL: <http://mpmath.org/>

5.12. MPFRGX

FUNCTIONAL DESCRIPTION

Mpfrgx is a library for the arithmetic of univariate polynomials over arbitrary precision real (Mpfr) or complex (Mpc) numbers, without control on the rounding. For the time being, only the few functions needed to implement the floating point approach to complex multiplication are implemented. On the other hand, these comprise asymptotically fast multiplication routines such as Toom-Cook and the FFT.

- Participant: Andreas Enge
- Contact: Andreas Enge
- URL: <http://www.multiprecision.org/index.php?prog=mpfrgx>

5.13. Nemo

FUNCTIONAL DESCRIPTION Nemo is a computer algebra package for the Julia programming language maintained by William Hart with code by William Hart, Tommy Hofmann, Claus Fieker, Fredrik Johansson, Oleksandr Motsak).

The features of Nemo include multiprecision integers and rationals, integers modulo n , p -adic numbers, finite fields (prime and non-prime order), number field arithmetic, maximal orders of number fields, arithmetic of ideals in maximal orders, arbitrary precision real and complex balls, generic polynomials, power series, fraction fields, residue rings and matrices.

- Participant: Fredrik Johansson
- Contact: William Hart
- URL: <http://nemocas.org/>

5.14. PARI/GP

FUNCTIONAL DESCRIPTION

Pari/Gp is a widely used computer algebra system designed for fast computations in number theory (factorisation, algebraic number theory, elliptic curves, ...), but it also contains a large number of other useful functions to compute with mathematical entities such as matrices, polynomials, power series, algebraic numbers, etc., and many transcendental functions.

- Participants: Karim Belabas, Bill Allombert, Henri Cohen and Andreas Enge
- Contact: Karim Belabas
- URL: <http://pari.math.u-bordeaux.fr/>

MAGIQUE-3D Project-Team

5. New Software and Platforms

5.1. Elasticus

SCIENTIFIC DESCRIPTION Elasticus simulate acoustic and elastic wave propagation in 2D and in 3D, using Discontinuous Galerkin Methods. The space discretization is based on two kind of basis functions, using Lagrange or Jacobi polynomials. Different kinds of fluxes (upwind and centered) are implemented, coupled with RK2 and RK4 time schemes.

FUNCTIONAL DESCRIPTION Elasticus is a sequential library, independent of DIVA and developed in Fortran, to simulate wave propagation in geophysical environment, based on a DG method. It is meant to help PhD students and post-doctoral fellows to easily implement their algorithms in the library. Thus, readability of the code is privileged to optimization of its performances. Developed features should be easily transferred in the computing platform of Total. Contrary to DIVA which only computes approximate solutions with P1, P2 and P3 elements, Elasticus manages arbitrary orders for the spatial discretization with DG method.

- Participants: Simon Ettouati, Julien Diaz and Lionel Boillot
- Partner: TOTAL
- Contact: Julien Diaz

5.2. Hou10ni

SCIENTIFIC DESCRIPTION Hou10ni simulates acoustic and elastic wave propagation in time domain and in harmonic domain, in 2D and in 3D. It is also able to model elasto-acoustic coupling. It is based on the second order formulation of the wave equation and the space discretization is achieved using Interior Penalty Discontinuous Galerkin Method. Recently, the harmonic domain solver has been extended to handle Hybridizable Discontinuous Galerkin Methods.

FUNCTIONAL DESCRIPTION This software simulates the propagation of waves in heterogeneous 2D and 3D media in time-domain and in frequency domain. It is based on an Interior Penalty Discontinuous Galerkin Method (IPDGM) and allows for the use of meshes composed of cells of various order (p-adaptivity in space).

- Participants: Julien Diaz, Marie Bonnasse-Gahot and Lionel Boillot
- Contact: Julien Diaz

5.3. Montjoie

SCIENTIFIC DESCRIPTION

Montjoie is designed for the efficient solution of time-domain and time-harmonic linear partial differential equations using high-order finite element methods. This code is mainly written for quadrilateral/hexahedral finite elements, partial implementations of triangular/tetrahedral elements are provided. The equations solved by this code, come from the "wave propagation" problems, particularly acoustic, electromagnetic, aeroacoustic, elasto-dynamic problems.

FUNCTIONAL DESCRIPTION

Montjoie is a code that provides a C++ framework for solving partial differential equations on unstructured meshes with finite element-like methods (continuous finite element, discontinuous Galerkin formulation, edge elements and facet elements). The handling of mixed elements (tetrahedra, prisms, pyramids and hexahedra) has been implemented for these different types of finite elements methods. For time-domain simulations, a wide range of ODE (Ordinary Differential Equation) solvers have been implemented : high-order explicit or implicit time schemes. Several applications are currently available : wave equation, elastodynamics, aero-acoustics, Maxwell's equations.

- Participants: Marc Duruflé, Juliette Chabassier, Mamadou N'Diaye and Michaël Leguèbe
- Contact: Marc Duruflé
- URL: <http://montjoie.gforge.inria.fr/>

5.4. TMBM-DG

Time-Marching Based Methods-Discontinuous Galerkin

SCIENTIFIC DESCRIPTION TMBM-DG simulate acoustic and elastic wave propagation in 2D and in 3D, using Discontinuous Galerkin Methods. The space discretization is based on two kind of basis functions, using Lagrange or Jacobi polynomials. Different kinds of fluxes (upwind and centered) are implemented, coupled with RK2 and RK4 time schemes.

FUNCTIONAL DESCRIPTION TMBM-DG is the follow up to DIVA-DG that we develop in collaboration with our partner Total. Its purpose is more general than DIVA-DG and should contains various DG schemes, basis functions and time schemes. It models wave propagation in acoustic media, elastic (isotropic and TTI) media and elasto-acoustic media, in two and three dimensions.

- Participants: Julien Diaz, Lionel Boillot and Simon Ettouati
- Partner: TOTAL
- Contact: Julien Diaz

MANAO Project-Team

6. New Software and Platforms

6.1. ALTA Lib

The ALTA Library

KEYWORDS: Statistic analysis - Fitting - Measures

FUNCTIONAL DESCRIPTION

ALTA is a multi-platform software library to analyze, fit and understand BRDFs. It provides a set of command line software to fit measured data to analytical forms, tools to understand models and data.

- Participants: Laurent Belcour, Romain Pacanowski, Xavier Granier and Pascal Barla
- Partner: LP2N (CNRS - UMR 5298)
- Contact: Romain Pacanowski
- URL: <http://alta.gforge.inria.fr/>

6.2. Elasticity Skinning

SCIENTIFIC DESCRIPTION

Geometric skinning techniques are very popular in the industry for their high performances, but fail to mimic realistic deformations. With elastic implicit skinning the skin stretches automatically (without skinning weights) and the vertices distribution is more pleasing. Our approach is more robust, for instance the angle's range of joints is larger than implicit skinning.

This software has been ported as a plugin for the Modo software (The Foundry) in collaboration with *Toulouse Tech Transfer*. This plugin has been bought by The Foundry, which maintains and sells it.

- Participants: Rodolphe Vaillant, Loïc Barthe, Florian Canezin, Gaël Guennebaud, Marie-Paule Cani, Damien Rohmer, Brian Wyvill, Olivier Gourmel and Mathias Paulin
- Partners: Université de Bordeaux - CNRS - INP Bordeaux - Université de Toulouse - Institut Polytechnique de Grenoble - Ecole Supérieure de Chimie Physique Electronique de Lyon
- Contact: Gaël Guennebaud
- URL: <http://rodolphe-vaillant.fr/?e=59>

6.3. Eigen

FUNCTIONAL DESCRIPTION

Eigen is an efficient and versatile C++ mathematical template library for linear algebra and related algorithms. In particular it provides fixed and dynamic size matrices and vectors, matrix decompositions (LU, LLT, LDLT, QR, eigenvalues, etc.), sparse matrices with iterative and direct solvers, some basic geometry features (transformations, quaternions, axis-angles, Euler angles, hyperplanes, lines, etc.), some non-linear solvers, automatic differentiations, etc. Thanks to expression templates, Eigen provides a very powerful and easy to use API. Explicit vectorization is performed for the SSE, AVX, FMA, AVX512, Altivec, VSX and ARM NEON instruction sets, with graceful fallback to non-vectorized code. Expression templates allow to perform global expression optimizations, and to remove unnecessary temporary objects.

In 2016, we released three revisions of the 3.2 branch, as well as the new 3.3 version that leverages numerous major novel features and improvements. Those include, a novel evaluation mechanism of expressions, support for AVX, FMA, AVX512, VSX and ZVector vector instructions, unaligned vectorization, nvcc/CUDA, more OpenMP parallelism, a fast divide and conquer SVD algorithm, a CompleteOrthogonalDecomposition class for fast minimal norm solving, a LS-CG solver, a fast reciprocal condition number estimators in LU and Cholesky factorizations, LU::transpose()/adjoint() API, support for inplace decompositions, support for matrix-free iterative solvers, new array functions, support for any BLAS/LAPACK libraries as backend, improved support for mixing scalar types, eigenvectors in GeneralizedEigenSolver, a complete rewrite of LinSpaced, a non officially supported but massively used Tensor module with CUDA and OpenCL support, and **more**.

- Participant: Gaël Guennebaud
- Contact: Gaël Guennebaud
- URL: <http://eigen.tuxfamily.org/>

6.4. HDRSee

KEYWORDS: OpenGL-GLSL HDR/LDR Viewer **FUNCTIONAL DESCRIPTION** HDRSee is a OpenGL/GLSL software that displays High Dynamic Range (HDR) and Low Dynamic Range (LDR) images. It is based on several libraries (e.g., glut, see below for full dependencies). To display HDR images, HDRSee implements a few tone-mapping operators. Moreover, it is designed with a plugin mechanism that let developers add, as easily as possible, their own tone-mapping operator. All tone-mapping operations are done using Graphics Hardware through pixel shader operations. The GUI currently used is nvWidgets.

- Participants: Romain Pacanowski, Xavier Granier.
- Partner: LP2N (CNRS - UMR 5298)
- Contact: Romain Pacanowski
- URL: <http://mhdrviewer.gforge.inria.fr/>

6.5. Patate Lib

KEYWORDS: Expressive rendering - Multi-scale analysis - Material appearance - Vector graphics - 2D animation

FUNCTIONAL DESCRIPTION

Patate is a header only C++/CUDA library for graphics applications. It provides a collection of Computer Graphics techniques that incorporate the latest innovations from Inria research teams working in the field. It strives for efficiency and ease-of-use by focusing on low-level core operators and key algorithms, organized in modules, each tackling a specific set of issues. The central goal of the library is to drastically reduce the time and efforts required to turn a research paper into a ready-to-use solution, for both commercial and academic purposes.

The library is still in its infancy and we are actively working on it to include the latest of our published research techniques. Modules will be dealing with graphics domains as varied as multi-scale analysis, material appearance, vector graphics, expressive rendering and 2D animation.

- Participants: Gaël Guennebaud, Pascal Barla, Simon Boyé, Gautier Ciaudo and Nicolas Mellado
- Contact: Gaël Guennebaud
- URL: <http://patate.gforge.inria.fr/html/>

MEMPHIS Project-Team

6. New Software and Platforms

6.1. COCOFLOW

FUNCTIONAL DESCRIPTION

The code is written in fortran 95 with a MPI parallelization. It solves equations of conservation modeling 3D compressible flows with elastic models as equation of state.

- Contact: Florian Bernard
- URL: <https://gforge.inria.fr/projects/cocoflow>

6.2. KOPPA

Kinetic Octree Parallel PolyAtomic

FUNCTIONAL DESCRIPTION

KOPPA is a C++/MPI numerical code solving a large range of rarefied flows from external to internal flows in 1D, 2D or 3D. Different kind of geometries can be treated such as moving geometries coming from CAO files or analytical geometries. The models can be solved on Octree grids with dynamic refinement.

- Participant: Florian Bernard
- Contact: Florian Bernard
- URL: <https://git.math.cnrs.fr/gitweb/?p=plm/fbernard/KOPPA.git;a=summary>

6.3. NS-penal

Navier-Stokes-penalization

KEYWORDS: 3D - Incompressible flows - 2D

FUNCTIONAL DESCRIPTION

The software can be used as a black box with the help of a data file if the obstacle is already proposed. For new geometries the user has to define them. It can be used with several boundary conditions (Dirichlet, Neumann, periodic) and for a wide range of Reynolds numbers.

- Partner: Université de Bordeaux
- Contact: Charles-Henri Bruneau

6.4. NaSCar

Navier-Stokes Cartesian

KEYWORDS: HPC - Numerical analyse - Fluid mechanics - Langage C - PETSc

SCIENTIFIC DESCRIPTION

NaSCar can be used to simulate both hydrodynamic bio-locomotion as fish like swimming and aerodynamic flows such wake generated by a wind turbine.

FUNCTIONAL DESCRIPTION

This code is devoted to solve 3D-flows in around moving and deformable bodies. The incompressible Navier-Stokes equations are solved on fixed grids, and the bodies are taken into account thanks to penalization and/or immersed boundary methods. The interface between the fluid and the bodies is tracked with a level set function or in a Lagrangian way. The numerical code is fully second order (time and space). The numerical method is based on projection schemes of Chorin-Temam's type. The code is written in C language and use Petsc library for the resolution of large linear systems in parallel.

NaSCar can be used to simulate both hydrodynamic bio-locomotion as fish like swimming and aerodynamic flows such wake generated by a wind turbine.

- Participant: Michel Bergmann
- Contact: Michel Bergmann
- URL: <https://gforge.inria.fr/projects/nascar/>

MNEMOSYNE Project-Team

6. New Software and Platforms

6.1. Positioning

Our previous works in the domain of well-defined distributed asynchronous adaptive computations [55], [52], [57] have already made us define a library (DANA [51]), closely related to both the notion of artificial neural networks and cellular automata. From a conceptual point of view, the computational paradigm supporting the library is grounded on the notion of a unit that is essentially a (vector of) potential that can vary along time under the influence of other units and learning. Those units can be organized into layers, maps and networks.

We will also have to interact with the High Performance Computing (HPC) community, since having large scale simulations at that mesoscopic level is an important challenge in our systemic view of computational neuroscience. Our approach implies to emulate the dynamics of thousands, or even millions, of integrated computational units, each of them playing the role of a whole elementary neural circuit (e.g. the microcolumn for the cortex). Mesoscopic models are considered in such an integrative approach, in order to exhibit global dynamical effect that would be hardly reachable by compartment models involving membrane equations or even spiking neuron networks.

The vast majority of high performance computing softwares for computational neuroscience addresses sub-neural or neural models [41], but coarser grained population models are also demanding for large scale simulations, with fully distributed computations, without global memory or time reference, as it is specified in (*cf.* § 3.2).

6.2. DANA

Distributed Asynchronous Numerical & Adaptive computing framework
FUNCTIONAL DESCRIPTION

DANA is a python framework whose computational paradigm is grounded on the notion of a unit that is essentially a set of time dependent values varying under the influence of other units via adaptive weighted connections. The evolutions of a unit's value are defined by a set of differential equations expressed in standard mathematical notation which greatly ease their definition. The units are organized into groups that form a model. Each unit can be connected to any other unit (including itself) using a weighted connection. The DANA framework offers a set of core objects needed to design and run such models. The modeler only has to define the equations of a unit as well as the equations governing the training of the connections. The simulation is completely transparent to the modeler and is handled by DANA. This allows DANA to be used for a wide range of numerical and distributed models as long as they fit the proposed framework (e.g. cellular automata, reaction-diffusion system, decentralized neural networks, recurrent neural networks, kernel-based image processing, etc.).

- Participant: Nicolas Rougier
- Contact: Nicolas Rougier
- URL: <http://dana.loria.fr/>

6.3. ENAS

Event Neural Assembly Simulation
KEYWORDS: Neurosciences - Health - Physiology
SCIENTIFIC DESCRIPTION

As one gains more intuitions and results on the importance of concerted activity in spike trains, models are developed to extract potential canonical principles underlying spike coding. These methods shed a new light on spike train dynamics. However, they require time and expertise to be implemented efficiently, making them hard to use in a daily basis by neuroscientists or modelers. To bridge this gap, we developed the license free multiplatform software ENAS (<https://enas.inria.fr>) integrating tools for individual and collective spike analysis and simulation, with some specificities devoted to the retina. The core of ENAS is the statistical analysis of population codes. One of its main strength is to provide statistical analysis of spike trains using Maximum Entropy-Gibbs distributions taking into account both spatial and temporal correlations as constraints, allowing to introduce causality and memory in statistics. ENAS also generates simulated spike trains. On one hand, one can draw a population raster from an user-specified Gibbs distribution. On the other hand, we have integrated in ENAS our retina simulator VIRTUAL RETINA, extended here to include lateral connections in the IPL. We hope that ENAS will become a useful tool for neuroscientists to analyse spike trains and we hope to improve it thanks to user feedback. Our goal is to progressively enrich it with the latest research results, in order to facilitate transfer of new methods to the community.

FUNCTIONAL DESCRIPTION

As one gains more intuitions and results on the importance of concerted activity in spike trains, models are developed to extract potential canonical principles underlying spike coding. These methods shed a new light on spike train dynamics. However, they require time and expertise to be implemented efficiently, making them hard to use in a daily basis by neuroscientists or modelers. To bridge this gap, we developed the license free multiplatform software ENAS integrating tools for spike trains analysis and simulation. These tools are accessible through a friendly Graphical User Interface that avoids any scripting or writing code from the user. Most of them have been implemented to run in parallel to reduce the time and memory consumption. ENAS offers basic visualizations and classical analysis for statistics of spike trains analysis. It also proposes statistical analysis with Maximum Entropy-Gibbs distributions taking into account both spatial and temporal correlations as constraints, allowing to introduce causality and memory in statistics. ENAS also includes specific tools dedicated to the retina: Receptive Field computation and a virtual retina simulator. Finally, ENAS generates synthetic rasters, either from know statistics or from the VIRTUAL RETINA simulator. We expect ENAS to become a useful tool for neuroscientists to analyse spike trains and we hope to improve it thanks to users feedback. From our perspective, our goal is to progressively enrich ENAS with the latest research results, in order to facilitate transfer of new methods to the community.

- Participants: Bruno Cessac, Salim Kraria, Hassan Nasser, Thierry Viéville, Rodrigo Cofre Torres, Geoffrey Portelli, Pierre Kornprobst, Theodora Karvouniari and Daniela Pamplona
- Contact: Bruno Cessac
- URL: <https://enas.inria.fr>

6.4. Virtual Enaction

KEYWORDS: Neurosciences - Simulation - Health

FUNCTIONAL DESCRIPTION

VirtualEnaction: A Platform for Systemic Neuroscience Simulation. The computational models studied in our team have applications that extend far beyond what is possible to experiment yet in human or non-human primate subjects. Real robotics experimentations are also impaired by rather heavy technological constraints, for instance, it is not easy to dismantle a given embedded system in the course of emerging ideas. The only versatile environment in which such complex behaviors can be studied both globally and at the level of details of the available modeling is a virtual environment, as in video games, Such a system can be implemented as “brainy-bot” (a programmed player based on our knowledge of the brain architecture) which goal is to survive in a complete manipulable environment.

In order to attain this rather ambitious objective we have deployed an existing open-source video game middleware (Minecraft) in order to be able to shape the survival situation to be studied and we have begun to revisit some models in order to be able to integrate them as an effective brainy-bot. This was made as a platform associated to a scenario that is the closest possible to a survival situation (foraging, predator-prey relationship, partner approach to reproduction). We could integrate in the platform an artificial agent with sensory inputs (visual, touch and smell), emotional and somatosensory cues (hunger, thirst, fear, ..) and motor outputs (movement, gesture, ..) connected to a "brain" whose architecture was corresponding to the major anatomical regions studied in the team.

Nevertheless, we have seen recently that a major entertainment company (Microsoft) bought Minecraft to make similar (but larger) adaptations to what was being targeted by our VirtualEnaction project. We are currently studying the resulting product (the Malmö project) in order to adapt our strategy.

- Participants: André Garenne, Frédéric Alexandre, Nicolas Rougier and Thierry Viéville
- Contact: Frédéric Alexandre
- URL: <http://virtualenaction.gforge.inria.fr/>

MONC Project-Team

6. New Software and Platforms

6.1. CADMOS

KEYWORDS: Health - Cancer - Partial differential equation - Cartesian grid

- Participants: Olivier Saut, Julien Jouganous, Annabelle Collin and Olivier Gallinato
- Partners: CNRS - INP Bordeaux - Université de Bordeaux
- Contact: Olivier Saut
- URL: <https://team.inria.fr/monc/software/>

6.2. Carcinom

Computer-Assisted Research about Cancer growth and INsights on Oncological Mechanisms

KEYWORDS: Cancer - Data modeling - Regression

- Participants: Vivien Pianet and Simon Evain
- Contact: Sébastien Benzekry
- URL: <https://team.inria.fr/monc/software/>

6.3. MetaPoumon

KEYWORDS: Health - Evolution - Cancer - Medical imaging

FUNCTIONAL DESCRIPTION

The software evaluates the aggressiveness of pulmonary metastasis or response to treatment for predictive goal. To do this, we use a mathematical model based on a set of equations to nonlinear partial differential equations. This model is calibrated to the patient data using a longitudinal sequence of CT or MRI of the patient.

- Participants: Olivier Saut, Thierry Colin, Marie Martin and Julien Jouganous
- Partners: CNRS - IPB - Université de Bordeaux
- Contact: Olivier Saut
- URL: <https://team.inria.fr/monc/software/>

6.4. Nenuphar

KEYWORDS: Modeling - Oncologie - Cancer - Partial differential equation - Medical - Medical imaging

FUNCTIONAL DESCRIPTION

The goal of project is to evaluate the aggressiveness of a tumor or its response to therapy. For that purpose, we use a mathematical model based on a set of nonlinear partial differential equations. This model is calibrated on patient data using a longitudinal sequence of CT Scan or MRI of the patient. This approach has been validated on about 35 clinical cases of lung metastases from various primary tumors (kidney, bladder, thyroid). Using two initial images showing the targeted lesion, we recover the patient-specific parameters of the model. The evolution of the disease is then predicted by letting the model run for later times with these parameters.

- Partners: CNRS - INP Bordeaux - Université Bordeaux I
- Contact: Marie Martin
- URL: <https://team.inria.fr/monc/software/>

6.5. PapriK

- Contact: Cynthia Perier
- URL: <https://team.inria.fr/monc/software/>

6.6. SESAR

Monitor of the effect of RT on Retroperitoneal Sarcoma

KEYWORDS: Segmentation - Health - DICOM - Cancer - Medical imaging

- Partner: Institut Bergonié
- Contact: Cynthia Perier
- URL: <https://team.inria.fr/monc/software/>

6.7. SegmentIt

KEYWORDS: Health - Signal - Registration of 2D and 3D multimodal images - 3D - Image analysis - Image - Processing - Medical imaging

FUNCTIONAL DESCRIPTION

Image processing software for anatomical and functional data. Segmentation, registration and digital filtering. Assesment of the kidney perfusion and the kidney function (to be continued).

- Participants: Thierry Colin, Olivier Saut, Vivien Pianet, Agathe Peretti, Marie Martin, Sébastien Benzekry, Baudoin Denis De Senneville, Cynthia Perier, Benjamin Taton, Nicolas Grenier and Christian Combe
- Contact: Benjamin Taton
- URL: <https://team.inria.fr/monc/software/>

PHOENIX Project-Team

6. New Software and Platforms

6.1. College +

KEYWORDS: Neurosciences - Health - Autism - Mobile application

School+ (or College+ in french) is a package of 7 applications. Three applications are assistive applications, guiding the child doing specific tasks. Three others are training applications made as serious games, addressing specific skills. The last application is a meta-application, comprising a link to the three training applications, with an access to statistics of their usage. For each application, data are separated from the design, meaning that every element of each application (pictures, texts, settings, etc.) can be changed at any time. Each application records a log file containing all the interactions performed by the child.

6.1.1. Assistive applications




Figure 2. Assistive applications

6.1.1.1. Routines application

This application shows a list of tasks, with a short description. After clicking the starting button, a specific slideshow is shown; it decomposes a task into steps. For each step, a text and a picture can be displayed. Thumbnail of previous and next steps are also displayed. This application guides the child through classroom situations: entering classroom, taking school materials out of a backpack, writing notes, handling agenda, leaving the classroom.

6.1.1.2. Communication application

With the same design, the assistance provided by this application targets to communicating situations inside the classroom. The application covers four scenarios addressing two interaction situations (initiating and answering the interaction) and two types of interlocutors (professor and classmate). For each scenario, different slideshows guide the child, depending on the goal of the interaction.

6.1.1.3. Emotion Regulation application

This application aims to assist the child to self-regulate his/her emotions. Four simplified emoticons are proposed to the child to choose from: anger, sadness, joy and fear. Then, (s)he selects a level of intensity via a thermometer with a scale from 1 to 4. In response, the application delivers different multimedia contents according to the level selected to help the child regulate his/her emotions. Typically, a text (breathing instructions) are shown at level 1, pictures at level 2, a video at level 3 and another text at level 4.

6.1.2. Training applications

These three applications are serious games with increasing levels of difficulties, reachable after a ratio of good answers has been attained.




Figure 3. Training applications

6.1.2.1. Emotion Recognition application with pictures

In this application, the child is instructed to identify a specific emotion among 4 pictures showing different people exhibiting an emotion. Seven emotions are involved in this application: joy, sadness, fear, anger, surprise, disgust and neutral. The emotion to be recognized is displayed together with its simplified emoticon. The type of pictures changes with the difficulty level: level 1 contains pictures of unfamiliar people and level 2 contains pictures of friends and relatives of the child.

6.1.2.2. Emotion Recognition application with videos

In this application, the child is presented with a fragment of an animated cartoon. At some point, the video stops and the child is asked to identify the emotion of the character. Four emotions are involved in this application: joy, sadness, fear and anger. Videos are slowed down, with a speed percentage that can be changed at each level. Videos change with difficulty level: level 1 contains videos of a very basic cartoon (only one cartoon character drawn by basic form un-textured), level 2 contains a video of more sophisticated cartoons and level 3 contains movies with actors.

6.1.2.3. Attention Training

In this application, the child is presented a picture of a face and asked to make eye contact with it. Second, a symbol appears briefly in the eyes of the character. Third, the child is asked to identify the symbol shown in the previously displayed picture, to make sure he kept eye contact. The speed at which the symbol appears and disappears is changed according to the difficulty level. Types of pictures also change with the level : level 1 contains pictures of faces and level 2 contains pictures of classroom situations.

- Participants: Damien Martin Guillerez, Charles Fage, Helene Sauzeon and Alexandre Spriet
- Contact: Charles Consel

6.2. DiaSuite

SCIENTIFIC DESCRIPTION

DiaSuite is a suite of tools covering the development life-cycle of a pervasive computing application:

6.2.1. Defining an application area

First, an expert defines a catalog of entities, whether hardware or software, that are specific to a target area. These entities serve as building blocks to develop applications in this area. They are gathered in a taxonomy definition, written in the taxonomy layer of the DiaSpec language.

6.2.2. Designing an application

Given a taxonomy, the architect can design and structure applications. To do so, the DiaSpec language provides an application design layer. This layer is dedicated to an architectural pattern commonly used in the pervasive computing domain. Describing the architecture application allows to further model a pervasive computing system, making explicit its functional decomposition.

6.2.3. Implementing an application

We leverage the taxonomy definition and the architecture description to provide dedicated support to both the entity and the application developers. This support takes the form of a Java programming framework, generated by the DiaGen compiler. The generated programming framework precisely guides the developer with respect to the taxonomy definition and the architecture description. It consists of high-level operations to discover entities and interact with both entities and application components. In doing so, it abstracts away from the underlying distributed technologies, providing further separation of concerns.

6.2.4. Testing an application

DiaGen generates a simulation support to test pervasive computing applications before their actual deployment. An application is simulated in the DiaSim tool, without requiring any code modification. DiaSim provides an editor to define simulation scenarios and a 2D-renderer to monitor the simulated application. Furthermore, simulated and actual entities can be mixed. This hybrid simulation enables an application to migrate incrementally to an actual environment.

6.2.5. Deploying a system

Finally, the system administrator deploys the pervasive computing system. To this end, a distributed systems technology is selected. We have developed a back-end that currently targets the following technologies: Web Services, RMI, SIP and OSGI. This targeting is transparent for the application code. The variety of these target technologies demonstrates that our development approach separates concerns into well-defined layers. This development cycle is summarized in the Figure 2 .

FUNCTIONAL DESCRIPTION

DiaSuite is developed as a research project by the Inria/LaBRI Phoenix research group. The DiaSuite approach covers the development life-cycle of a pervasive computing application. It takes the form of a methodology, supported by (1) a high-level design language and (2) a suite of tools covering the development life-cycle of a pervasive computing application. Specifically, we have developed a design language dedicated to describing pervasive computing systems and a suite of tools providing customized support for each development stage of a pervasive computing system, namely, implementation (e.g., programming support), testing (e.g., unit test, 2D simulator), and deployment (e.g., distribution platforms like SIP and Web Services).

- Participants: Charles Consel, Milan Kabac, Paul Van Der Walt, Adrien Carteron and Alexandre Spriet
- Contact: Charles Consel

6.3. DiaSuiteBOX

KEYWORDS: Health - Smart home - Open application store - Development tool suite - Application certification - Home care

FUNCTIONAL DESCRIPTION




Figure 4. DiaSuite Development Cycle

DiaSuiteBOX proposes an application store that gathers the devices deployed at home. This store is open and available online such as an application store for Smartphone.

- Participants: Bertran Benjamin, Bruneau Julien, Consel Charles, Quentin Enard, Milan Kabac, Damien Martin Guillerez, Emilie Balland, Damien Cassou, Amelie Marzin, Julien Durand, Quentin Barlas, Ludovic Fornasari, Joan Rieu, Adrien Carteron, Eugene Volanschi and Helene Sauzeon
- Partners: CNRS - IPB - Université de Bordeaux
- Contact: Charles Consel
- URL: <https://diasuitebox.inria.fr/>

6.4. DomAssist

KEYWORDS: Health - Mobile application - Persons attendant - Home care

The HomeAssist platform (or DomAssist in french) proposes a systemic approach to introducing an assistive technological platform for older people. To do so, we formed a trans-disciplinary team that allows (1) to identify the user needs from a gerontological and psychological viewpoint; (2) to propose assistive applications designed by human factors and HCI experts, in collaboration with caregivers and users; (3) to develop and test applications by software engineers; (4) to conduct a field study for assessing the benefits of the platform and assistive applications, in collaboration with caregivers, by deploying the system at the actual home of older adults.

The HomeAssist platform is implemented on top of the DiaSuiteBox platform, using a suite of tools, namely DiaSuite, that have been designed, developed and tested by our research group at Inria. The DiaSuite tools include a dedicated integrated development environment that enables applications to be developed quickly and safely. This technology has been successfully applied to a variety of domains where environments consist of networked objects that need to be orchestrated.

6.4.1. Applications

HomeAssist offers an online catalog of applications. Using this catalog, the user and the caregiver determine what and how activities should be assisted by selecting the appropriate assistive applications and configuring them with respect to the user’s requirements and preferences. The resulting set of applications forms a personalized assistive support. Additionally, to respond to evolving needs, our platform allows to stop/remove applications easily and to install new ones from the online catalog.

This platform proposes many applications in three domains of everyday life.

Daily activities: including activity monitoring, light path, and a reminder.

Home or personal safety: including entrance monitoring, stove monitoring, and warning if no movements are detected after a certain amount of time.

Communications and social activities: including collaborative games, videoconference, information about local events, TV programming, etc.




Figure 5. The HomeAssist platform and applications

For video presentations of HomeAssist, see the following:

- <http://videotheque.inria.fr/videotheque/media/23705>. Title: “DiaSuiteBox”, 2013.
- <http://videotheque.inria.fr/videotheque/media/29998>. Title: “DomAssist : L’assistance numérique à la personne”, 2014.

6.4.2. Devices

Several entities have been identified to deliver an assistive support. These entities include (1) technological devices: wireless sensors (motion detectors, contact sensors and smart electric switches), and two tablets, and (2) software services (agenda, address book, mail agent, and photo agent) to monitor everyday activities and

propose assistive applications. Sensors are placed in relevant rooms in the house: kitchen, bedroom, bathroom, and around the entrance.



Figure 6. HomeAssist devices

FUNCTIONAL DESCRIPTION

3 mobile applications for assistive living : (1) DiAndroid : Interface for the main tablet with the DiaSuiteBox applications including those for the daily activities, the meetings scheduling, etc. and for home and personal safety; (2) Accueil : home screen restraining the use of a secondary tablet and offering communications and social activities applications with simplified communication means (ie. eMail), collaborative games, etc.; (3) eMail : mail client made for older people.

- Participants: Alexandre Spriet, Charles Consel, Helene Sauzeon and Julien Durand
- Partners: CNRS - IPB - Université de Bordeaux
- Contact: Charles Consel
- URL: <http://phoenix.inria.fr/research-projects/homeassist>

PLEIADE Team

6. New Software and Platforms

6.1. Magus

KEYWORDS: Bioinformatics - Genomic sequence - Knowledge database

FUNCTIONAL DESCRIPTION

Comparative genomics requires efficient and scalable tools for managing knowledge about genomes, genes, and the high-dimensional relations between them.

The MAGUS genome annotation system integrates genome sequences and sequences features, in silico analyses, and views of external data resources into a familiar user interface requiring only a Web navigator. MAGUS implements annotation workflows and enforces curation standards to guarantee consistency and integrity. As a novel feature the system provides a workflow for simultaneous annotation of related genomes through the use of protein families identified by in silico analyses this has resulted in a three-fold increase in curation speed, compared to one-at-a-time curation of individual genes. This allows us to maintain standards of high-quality manual annotation while efficiently using the time of volunteer curators.

MAGUS can be used on small installations with a web server and a relational database on a single machine, or scaled out in clusters or elastic clouds using Apache Cassandra for NoSQL data storage and Apache Hadoop for Map-Reduce.

- Participants: David Sherman, Pascal Durrens
- Partners: CNRS - INRA - Université de Bordeaux
- Contact: David James Sherman
- URL: <http://magus.gforge.inria.fr>

6.2. Mimoza

KEYWORDS: Systems Biology - Bioinformatics - Biotechnology

FUNCTIONAL DESCRIPTION

Mimoza uses metabolic model generalization and cartographic paradigms to allow human experts to explore a metabolic model in a hierarchical manner. The software creates a zoomable representation of a model submitted by the user in SBML format. The most general view represents the compartments of the model, the next view shows the visualization of generalized versions of reactions and metabolites in each compartment, and the most detailed view visualizes the initial model with the generalization-based layout (where similar metabolites and reactions are placed next to each other). The zoomable representation is implemented using the Leaflet JavaScript library for mobile-friendly interactive maps. Users can click on reactions and compounds to see the information about their annotations. The resulting map can be explored on-line, or downloaded in a COMBINE archive.

- Participants: Anna Zhukova and David James Sherman
- Contact: David James Sherman
- URL: <http://mimoza.bordeaux.inria.fr/>

6.3. Pantograph

KEYWORDS: Systems Biology - Bioinformatics - Genomics - Gene regulatory networks

FUNCTIONAL DESCRIPTION

Pantograph is a software toolbox to reconstruct, curate and validate genome-scale metabolic models. It uses existing metabolic models as templates, to start its reconstructions process, to which new, species-specific reactions are added. Pantograph uses an iterative approach to improve reconstructed models, facilitating manual curation and comparisons between reconstructed model's predictions and experimental evidence.

Pantograph uses a consensus procedure to infer relationships between metabolic models, based on several sources of orthology between genomes. This allows for a very detailed rewriting of reaction's genome associations between template models and the model you want to reconstruct.

- Participants: Nicolas Loira, Anna Zhukova, David James Sherman and Pascal Durrens
- Partner: University of Chile
- Contact: Nicolas Loira
- URL: <http://pathhtastic.gforge.inria.fr/>

6.4. BioRica

KEYWORDS: Systems Biology - Bioinformatics - Hierarchical models - Hybrid models - Stochastic models
FUNCTIONAL DESCRIPTION

BioRica is used to mathematically describe the behavior of complex biological systems.

It is a software platform that permits simulation of biological systems on the basis of their description. It allows one to reuse existing biological models and to combine them into more complex models.

- Partner: University of Chile
- Contact: David Sherman
- URL: <http://biorica.gforge.inria.fr/>

6.5. Declic

Metabarcoding relies on mapping large sets of reads on reliable databases, with taxonomically annotated sequences. Declic facilitates data analyses for metabarcoding.

FUNCTIONAL DESCRIPTION

Declic is a Python library that provides several tools for data analysis in the domains of multivariate data analysis, machine learning, and graph based methods. It can be used to study in-depth the accuracy of the dictionary between molecular based and morphological based taxonomy.

Declic includes an interpreter for a Domain Specific Language (DSL) to make its Python library easy to use for scientists familiar with environments such as R.

- Partner: INRA
- Contact: Alain Franc

6.6. Platforms

6.6.1. Plafrim

Plafrim (<http://plafrim.fr>) is an essential instrument for PLEIADE. We use it for developing software data analysis methods and evaluating them at real world scale. The platform combines considerable computing power with excellent support, both in terms of the quality of the interactions with the local staff and of the ease of large-scale data transfer between Plafrim and PLEIADE's data storage infrastructure. Plafrim facilitates collaboration between team members who are not in the Bordeaux Sud-Ouest building, and furthermore allows us to share best practices and tools with other teams from the Center.

6.6.2. Inria forge and Inria continuous integration

The Inria forge (<http://gforge.inria.fr>) provides a secure collaboration platform for software project administration and source code management, and Inria's continuous integration platform (<http://ci.inria.fr>) provides a cloud-based service for automatic compilation and testing of software systems. PLEIADE uses these two services extensively for agile software development. The continuous integration platform allows us to verify the correct operation of our methods in different operating system and deployment environments.

6.6.3. Team Platform

PLEIADE maintains a dedicated computing platform for software development and experimentation by the team, comprised of a private cloud, storage, and a Project Atomic cluster for hosting Docker containers

POSET Team

6. New Software and Platforms

6.1. T-calculus

Sketched in [10], the *T-calculus* is a Domain Specific Language⁰ to provide simple and robust high-level description mechanisms of reactive systems. It will offer a programmatic view of the tile modeling paradigm [3], [9]. Its definition has been refined a number of times (see e.g. [10], [8]). A prototype implementation of its reactive kernel has eventually been achieved in Haskell on top of the Euterpea libraries during the spring 2016 [15], [21], [19], [23]

⁰See [30] for an early note by Hudak about the notion of Domain Specific Language, and see [29], [32] for application of this notion is computer music.

POTIOC Project-Team

6. New Software and Platforms

6.1. PapARt

PapARt is a software development kit (SDK) that enables the creation of interactive projection mapping (See <https://project.inria.fr/papart>). This year, we focused on making this toolkit widely available. We created a set of examples and created tutorials. The PapARt code is now Open Source, our objective being to favor a wide appropriation by artists, teachers, or students.

Participants: Jeremy Laviolle, Martin Hachet

URL: <https://github.com/poqudrof/Papart-examples/wiki>

6.2. Helios

Helios is a software tool (Unity3D) we have developed in collaboration with Stéphanie Fleck from Université de Lorraine. It is dedicated to the learning of astronomy at school. It bases on augmented reality and tangible interaction. See Section 7.4 .

Participants: Robin Gourdel, Jérémy Laviolle, Benoit Coulais, Martin Hachet.

Partners: Université de Lorraine - SATT Nancy Grand-Est.

6.3. Aïana

We have developped Aïana, a MOOC player, with the support of the Inria MOOC Lab. Aïana offers original interaction features in order to enable a wide spectrum of users including persons with disabilities. The first version of Aïana has been used by the 3700 participants of the Digital Accessibility MOOC we have produced on the national MOOC platform FUN. See Section 7.9 .

Participants: Pierre-Antoine Cinquin, Pascal Guitton

Partners: LearningLab Inria

6.4. HOBIT

Along with the project HOBIT (see Section 7.1), we continue enhancing the platform that is dedicated to the simulation and augmentation of optics experiments. In particular, this year, we did a major evolution that consists in making the system reconfigurable. Various optical components be plugged in, and the simulation and augmentations are updated accordingly.

Participants: Benoit Coulais, David Furio, Martin Hachet.

Partners: Université de Bordeaux - IUT de Bordeaux, LaBRI, IMS, CELIA

<https://project.inria.fr/hobit>

REALOPT Project-Team

6. New Software and Platforms

6.1. BaPCod : A generic Branch-And-Price Code

KEYWORDS: Column Generation - Branch-and-Price - Branch-and-Cut - Mixed Integer Programming - Mathematical Optimization - Benders Decomposition - Dantzig-Wolfe Decomposition - Extended Formulation

FUNCTIONAL DESCRIPTION

BaPCod is a prototype code that solves Mixed Integer Programs (MIP) by application of reformulation and decomposition techniques. The reformulated problem is solved using a branch-and-price-and-cut (column generation) algorithm, Benders approaches, or network flow algorithms.

- Participants: Francois Vanderbeck, Ruslan Sadykov, Issam Tahiri, Romain Leguay, Artur Alves Pessoa, Boris Detienne, Franck Labat, François Clautiaux, Pierre Pesneau, Eduardo Uchoa Barboza, Michael Poss and Halil Sen
- Partners: CNRS - IPB - Universidade Federal Fluminense - Université de Bordeaux
- Contact: François Vanderbeck
- URL: <https://wiki.bordeaux.inria.fr/realopt/pmwiki.php/Project/BaPCod>

SISTM Project-Team

6. New Software and Platforms

6.1. New software

6.1.1. *clogitLasso*

Lasso Estimation of Conditional Logistic Regression Models for Case-Crossover and Matched Case-Control Studies

- KEYWORDS: Classification , Statistics , Cluster , Machine learning , Regression
- FUNCTIONAL DESCRIPTION Fit a sequence of conditional logistic regression models with lasso, for small to large sized samples.
- Contact: Marius Kwémou
- URL: <https://github.com/robingenuer/>

6.1.2. *COVVSURF*

Combination of Clustering Of Variables and Variable Selection Using Random Forests

- KEYWORDS: Classification , Statistics , Cluster , Machine learning , Regression
- FUNCTIONAL DESCRIPTION This package implements a two stage strategy, where first we use ClustOfVar package to perform a clustering of variables and second we use VSURF package to select features (i.e. synthetic variables built in the first step).
- Contact: Robin Genuer
- URL: <https://github.com/robingenuer/CoVVSURF>

6.1.3. *NPflow*

Bayesian Nonparametrics for Automatic Gating of Flow-Cytometry Data

- KEYWORDS: Bayesian estimation , Bioinformatics , Biostatistics
- FUNCTIONAL DESCRIPTION Dirichlet process mixture of multivariate normal, skew normal or skew t-distributions modeling oriented towards flow-cytometry data pre-processing applications.
- Contact: Boris Hejblum
- URL: <https://cran.r-project.org/web/packages/NPflow/>

6.1.4. *tcgsaseq*

Time-Course Gene Set Analysis for RNA-Seq Data

- KEYWORDS: Genomics , Biostatistics , Statistical modeling , RNA-seq , Gene Set Analysis
- FUNCTIONAL DESCRIPTION Gene set analysis of longitudinal RNA-seq data with variance component score test accounting for data heteroscedasticity through precision weights.
- Contact: Boris Hejblum
- URL: <https://cran.r-project.org/web/packages/tcgsaseq/>

6.1.5. *CD4 Shiny*

Reference curves for CD4 response to antiretroviral therapy in HiV infected patients

- KEYWORDS: HIV infection, antiretroviral therapy, cd4 response, reference curves, quantile regression
- FUNCTIONAL DESCRIPTION References curves for CD4 response to antiretroviral therapy in HIV infected patients derived from large cohorts and estimated according to known factors associated with the response to antiretroviral therapy.
- Contact: Rodolphe Thiébaud
- URL: <http://shiny.isped.u-bordeaux.fr/CD4refcurves/>

6.2. Older software still maintained by SISTM

6.2.1. *NIMROD*

Normal approximation Inference in Models with Random effects based on Ordinary Differential equations

- KEYWORDS: Biostatistics - Optimization
- FUNCTIONAL DESCRIPTION We have written a specific program called NIMROD for estimating parameter of ODE based population models.
- Contact: Rodolphe Thiebaut
- URL: <http://etudes.isped.u-bordeaux2.fr/BIOSTATISTIQUE/NIMROD/documentation/html/index.html>

6.2.2. *R2GUESS*

Graphical processing Unit Evolutionary Stochastic Search

- KEYWORDS: Bioinformatics - Biostatistics
- FUNCTIONAL DESCRIPTION R2GUESS package is a wrapper of the GUESS (Graphical processing Unit Evolutionary Stochastic Search) program. GUESS is a computationally optimised C++ implementation of a fully Bayesian variable selection approach that can analyse, in a genome-wide context, single and multiple responses in an integrated way. The program uses packages from the GNU Scientific Library (GSL) and offers the possibility to re-route computationally intensive linear algebra operations towards the Graphical Processing Unit (GPU) through the use of proprietary CULA-dense library.
- Contact: Rodolphe Thiebaut
- URL: <https://cran.r-project.org/web/packages/R2GUESS/index.html>

6.2.3. *TcGSA*

Time-course Gene Set Analysis

- KEYWORDS: Bioinformatics - Genomics
- FUNCTIONAL DESCRIPTION An R package for the gene set analysis of longitudinal gene expression data sets. Under development, and soon to be available on the CRAN website, this package implements a Time-course Gene Set Analysis method and provides useful plotting functions facilitating the interpretation of the results.
- Contact: Boris Hejblum
- URL: <https://cran.r-project.org/web/packages/TcGSA/index.html>

6.2.4. *VSURF*

Variable Selection Using Random Forests

- KEYWORD: Bioinformatics
- FUNCTIONAL DESCRIPTION An R package for Variable Selection Using Random Forests. Available on CRAN, this package performs an automatic (meaning completely data-driven) variable selection procedure. Originally designed to deal with high dimensional data, it can also be applied to standard datasets.
- Contact: Robin Genuer
- URL: <https://cran.r-project.org/web/packages/VSURF/index.html>

6.2.5. *marqLevAlg*

Function optimization (minimization or maximization)

- KEYWORDS: Optimization - Biostatistics
- FUNCTIONAL DESCRIPTION An R package for function optimization. Available on CRAN, this package performs a minimization of function based on the Marquardt-Levenberg algorithm. This package is really useful when the surface to optimize is non-strictly convex or far from a quadratic function. A new convergence criterion, the relative distance to maximum (RDM), allows the user to have a better confidence in the stopping points, other than basic algorithm stabilization.
- Contact: Melanie Prague
- URL: <https://cran.r-project.org/web/packages/marqLevAlg/index.html>

STORM Team

5. New Software and Platforms

5.1. Chameleon

KEYWORDS: HPC - Dense linear algebra - Task-based algorithm - Runtime system - Task scheduling

SCIENTIFIC DESCRIPTION

Chameleon is part of the MORSE (Matrices Over Runtime Systems @ Exascale) project. The overall objective is to develop robust linear algebra libraries relying on innovative runtime systems that can fully benefit from the potential of those future large-scale complex machines.

We expect advances in three directions based first on strong and closed interactions between the runtime and numerical linear algebra communities. This initial activity will then naturally expand to more focused but still joint research in both fields.

1. Fine interaction between linear algebra and runtime systems. On parallel machines, HPC applications need to take care of data movement and consistency, which can be either explicitly managed at the level of the application itself or delegated to a runtime system. We adopt the latter approach in order to better keep up with hardware trends whose complexity is growing exponentially. One major task in this project is to define a proper interface between HPC applications and runtime systems in order to maximize productivity and expressivity. As mentioned in the next section, a widely used approach consists in abstracting the application as a DAG that the runtime system is in charge of scheduling. Scheduling such a DAG over a set of heterogeneous processing units introduces a lot of new challenges, such as predicting accurately the execution time of each type of task over each kind of unit, minimizing data transfers between memory banks, performing data prefetching, etc. Expected advances: In a nutshell, a new runtime system API will be designed to allow applications to provide scheduling hints to the runtime system and to get real-time feedback about the consequences of scheduling decisions.

2. Runtime systems. A runtime environment is an intermediate layer between the system and the application. It provides low-level functionality not provided by the system (such as scheduling or management of the heterogeneity) and high-level features (such as performance portability). In the framework of this proposal, we will work on the scalability of runtime environment. To achieve scalability it is required to avoid all centralization. Here, the main problem is the scheduling of the tasks. In many task-based runtime environments the scheduler is centralized and becomes a bottleneck as soon as too many cores are involved. It is therefore required to distribute the scheduling decision or to compute a data distribution that impose the mapping of task using, for instance the so-called “owner-compute” rule. Expected advances: We will design runtime systems that enable an efficient and scalable use of thousands of distributed multicore nodes enhanced with accelerators.

3. Linear algebra. Because of its central position in HPC and of the well understood structure of its algorithms, dense linear algebra has often pioneered new challenges that HPC had to face. Again, dense linear algebra has been in the vanguard of the new era of petascale computing with the design of new algorithms that can efficiently run on a multicore node with GPU accelerators. These algorithms are called “communication-avoiding” since they have been redesigned to limit the amount of communication between processing units (and between the different levels of memory hierarchy). They are expressed through Direct Acyclic Graphs (DAG) of fine-grained tasks that are dynamically scheduled. Expected advances: First, we plan to investigate the impact of these principles in the case of sparse applications (whose algorithms are slightly more complicated but often rely on dense kernels). Furthermore, both in the dense and sparse cases, the scalability on thousands of nodes is still limited, new numerical approaches need to be found. We will specifically design sparse hybrid direct/iterative methods that represent a promising approach.

The overall goal of the MORSE associate team is to enable advanced numerical algorithms to be executed on a scalable unified runtime system for exploiting the full potential of future exascale machines.

FUNCTIONAL DESCRIPTION

Chameleon is a dense linear algebra software relying on sequential task-based algorithms where sub-tasks of the overall algorithms are submitted to a Runtime system. A Runtime system such as StarPU is able to manage automatically data transfers between not shared memory area (CPUs-GPUs, distributed nodes). This kind of implementation paradigm allows to design high performing linear algebra algorithms on very different type of architecture: laptop, many-core nodes, CPUs-GPUs, multiple nodes. For example, Chameleon is able to perform a Cholesky factorization (double-precision) at 80 TFlop/s on a dense matrix of order 400 000 (e.i. 4 min).

- Participants: Emmanuel Agullo, Mathieu Faverge, Cédric Castagnede and Florent Pruvost
- Partners: Innovative Computing Laboratory (ICL) - King Abdullha University of Science and Technology - University of Colorado Denver
- Contact: Emmanuel Agullo
- URL: <https://project.inria.fr/chameleon/>

5.2. KLANG-OMP

The KStar OpenMP Compiler

KEYWORDS: Compilers - Task scheduling - OpenMP - Source-to-source compiler - Data parallelism

FUNCTIONAL DESCRIPTION

The Klang-Omp compiler, now renamed KStar following the recommendation of the local experimentation and development service, is a source-to-source OpenMP compiler for languages C and C++. The KStar compiler translates OpenMP directives and constructs into API calls from the StarPU runtime system or the XKaapi runtime system. The KStar compiler is virtually fully compliant with OpenMP 3.0 constructs. The KStar compiler supports OpenMP 4.0 dependent tasks and accelerated targets.

- Participants: Olivier Aumage, Nathalie Furmento, Samuel Pitoiset and Samuel Thibault
- Contact: Olivier Aumage
- URL: <http://kstar.gforge.inria.fr/#!index.md>

5.3. KaStORS

The KaStORS OpenMP Benchmark Suite

KEYWORDS: Benchmarking - HPC - Task-based algorithm - Task scheduling - OpenMP - Data parallelism

FUNCTIONAL DESCRIPTION

The KaStORS benchmarks suite has been designed to evaluate implementations of the OpenMP dependent task paradigm, introduced as part of the OpenMP 4.0 specification.

- Participants: Olivier Aumage, François Broquedis, Pierrick Brunet, Nathalie Furmento, Thierry Gautier, Samuel Thibault and Philippe Virouleau
- Contact: Thierry Gautier
- URL: <http://kastors.gforge.inria.fr/#!index.md>

5.4. MORSE

- Contact: Emmanuel Agullo
- URL: <http://icl.cs.utk.edu/morse/>

5.5. AFF3CT

A Fast Forward Error Correction Tool (previously named P-Edge).

KEYWORDS: Code generation - Error Correction Code

FUNCTIONAL DESCRIPTION

The AFF3CT library joins genericity techniques together with code generation capabilities to enable implementing efficient and portable error correction codes. The genericity offered allows to easily experiment with a large panel of algorithmic variants.

- Previous name: P-Edge
- Authors: Adrien Cassagne, Olivier Aumage, Bertrand Le Gal, Camille Leroux and Denis Barthou
- Partner: IMS
- Contact: Adrien Cassagne
- URL: <https://aff3ct.github.io/>

5.6. StarPU

The StarPU Runtime System

KEYWORDS: HPC - Scheduling - GPU - Multicore - Performance

SCIENTIFIC DESCRIPTION

Traditional processors have reached architectural limits which heterogeneous multicore designs and hardware specialization (eg. coprocessors, accelerators, ...) intend to address. However, exploiting such machines introduces numerous challenging issues at all levels, ranging from programming models and compilers to the design of scalable hardware solutions. The design of efficient runtime systems for these architectures is a critical issue. StarPU typically makes it much easier for high performance libraries or compiler environments to exploit heterogeneous multicore machines possibly equipped with GPGPUs or Cell processors: rather than handling low-level issues, programmers may concentrate on algorithmic concerns. Portability is obtained by the means of a unified abstraction of the machine. StarPU offers a unified offloadable task abstraction named "codelet". Rather than rewriting the entire code, programmers can encapsulate existing functions within codelets. In case a codelet may run on heterogeneous architectures, it is possible to specify one function for each architectures (eg. one function for CUDA and one function for CPUs). StarPU takes care to schedule and execute those codelets as efficiently as possible over the entire machine. In order to relieve programmers from the burden of explicit data transfers, a high-level data management library enforces memory coherency over the machine: before a codelet starts (eg. on an accelerator), all its data are transparently made available on the compute resource. Given its expressive interface and portable scheduling policies, StarPU obtains portable performances by efficiently (and easily) using all computing resources at the same time. StarPU also takes advantage of the heterogeneous nature of a machine, for instance by using scheduling strategies based on auto-tuned performance models.

StarPU is a task programming library for hybrid architectures

The application provides algorithms and constraints: - CPU/GPU implementations of tasks - A graph of tasks, using either the StarPU's high level GCC plugin pragmas or StarPU's rich C API

StarPU handles run-time concerns - Task dependencies - Optimized heterogeneous scheduling - Optimized data transfers and replication between main memory and discrete memories - Optimized cluster communications

Rather than handling low-level scheduling and optimizing issues, programmers can concentrate on algorithmic concerns!

FUNCTIONAL DESCRIPTION

StarPU is a runtime system that offers support for heterogeneous multicore machines. While many efforts are devoted to design efficient computation kernels for those architectures (e.g. to implement BLAS kernels on GPUs), StarPU not only takes care of offloading such kernels (and implementing data coherency across the machine), but it also makes sure the kernels are executed as efficiently as possible.

- Participants: Cédric Augonnet, Samuel Thibault, Nathalie Furmento, Simon Archipoff, Jérôme Clet-Ortega, Nicolas Collin, Ludovic Courtes, Mehdi Juhour, Xavier Lacoste, Benoît Lize, Ludovic Stordeur, Cyril Roelandt, Corentin Salingue, Chiheb Sakka, Samuel Pitoiset, François Tessier, Pierre-André Wacrenier, Andra Hugo, Terry Cojean, Raymond Namyst, Olivier Aumage and Marc Sergent
- Contact: Olivier Aumage
- URL: <http://starpu.gforge.inria.fr/>

5.7. hwloc

Hardware Locality

KEYWORDS: HPC - Topology - Open MPI - Affinities - GPU - Multicore - NUMA - Locality

FUNCTIONAL DESCRIPTION

Hardware Locality (hwloc) is a library and set of tools aiming at discovering and exposing the topology of machines, including processors, cores, threads, shared caches, NUMA memory nodes and I/O devices. It builds a widely-portable abstraction of these resources and exposes it to applications so as to help them adapt their behavior to the hardware characteristics. They may consult the hierarchy of resources, their attributes, and bind task or memory on them.

hwloc targets many types of high-performance computing applications, from thread scheduling to placement of MPI processes. Most existing MPI implementations, several resource managers and task schedulers, and multiple other parallel libraries already use hwloc.

- Participants: Brice Goglin and Samuel Thibault
- Partners: AMD - Intel - Open MPI consortium
- Contact: Brice Goglin
- URL: <http://www.open-mpi.org/projects/hwloc/>

TADAAM Team

6. New Software and Platforms

6.1. NetLoc

Network Locality

FUNCTIONAL DESCRIPTION

NETLOC (Network Locality) is a library that extends HWLOC to network topology information by assembling HWLOC knowledge of server internals within graphs of inter-node fabrics such as Infiniband, Intel OmniPath or Cray networks. NETLOC builds a software representation of the entire cluster so as to help application properly place their tasks on the nodes. It may also help communication libraries optimize their strategies according to the wires and switches. NETLOC targets the same challenges as HWLOC but focuses on a wider spectrum by enabling cluster-wide solutions such as process placement. NETLOC is distributed within HWLOC releases starting with HWLOC 2.0.

- Participants: Cyril Bordage and Brice Goglin
- Contact: Brice Goglin
- URL: <http://www.open-mpi.org/projects/netloc/>

6.2. NewMadeleine

KEYWORDS: High-performance calculation - MPI communication

FUNCTIONAL DESCRIPTION

NewMadeleine is the fourth incarnation of the Madeleine communication library. The new architecture aims at enabling the use of a much wider range of communication flow optimization techniques. Its design is entirely modular: drivers and optimization strategies are dynamically loadable software components, allowing experimentations with multiple approaches or on multiple issues with regard to processing communication flows.

The optimizing scheduler SchedOpt targets applications with irregular, multi-flow communication schemes such as found in the increasingly common application conglomerates made of multiple programming environments and coupled pieces of code, for instance. SchedOpt itself is easily extensible through the concepts of optimization strategies (what to optimize for, what the optimization goal is) expressed in terms of tactics (how to optimize to reach the optimization goal). Tactics themselves are made of basic communication flows operations such as packet merging or reordering.

The communication library is fully multi-threaded through its close integration with PIOMan. It manages concurrent communication operations from multiple libraries and from multiple threads. Its MPI implementation Mad-MPI fully supports the MPI_THREAD_MULTIPLE multi-threading level.

- Participants: Alexandre Denis, Nathalie Furmento, Raymond Namyst and Clement Foyer
- Contact: Alexandre Denis
- URL: <http://pm2.gforge.inria.fr/newmadeleine/>

6.3. PaMPA

Parallel Mesh Partitioning and Adaptation

KEYWORDS: Dynamic load balancing - Unstructured heterogeneous meshes - Parallel remeshing - Subdomain decomposition - Parallel numerical solvers

SCIENTIFIC DESCRIPTION

PAMPA is a parallel library for handling, redistributing and remeshing unstructured meshes on distributed-memory architectures. PAMPA dramatically eases and speeds-up the development of parallel numerical solvers for compact schemes. It provides solver writers with a distributed mesh abstraction and an API to:

- describe unstructured and possibly heterogeneous meshes, on the form of a graph of interconnected entities of different kinds (e.g. elements, faces, edges, nodes);
- attach values to the mesh entities;
- distribute such meshes across processing elements, with an overlap of variable width;
- perform synchronous or asynchronous data exchanges of values across processing elements;
- describe numerical schemes by means of iterators over mesh entities and their connected neighbors of a given kind;
- redistribute meshes so as to balance computational load;
- perform parallel dynamic remeshing, by applying adequately a user-provided sequential remesher to relevant areas of the distributed mesh.

PAMPA runs concurrently multiple sequential remeshing tasks to perform dynamic parallel remeshing and redistribution of very large unstructured meshes. E.g., it can remesh a tetrahedral mesh from 43 million elements to more than 1 billion elements on 280 Broadwell processors in 20 minutes.

FUNCTIONAL DESCRIPTION

Parallel library for handling, redistributing and remeshing unstructured, heterogeneous meshes on distributed-memory architectures. PAMPA dramatically eases and speeds-up the development of parallel numerical solvers for compact schemes.

- Participants: Cedric Lachat, François Pellegrini and Cécile Dobrzynski
- Partners: CNRS - IPB - Université de Bordeaux
- Contact: Cedric Lachat
- URL: <http://project.inria.fr/pampa/>

6.4. SCOTCH

KEYWORDS: High-performance computing - Graph algorithms - Domain decomposition - Static mapping - Mesh partitioning - Sparse matrix ordering

SCIENTIFIC DESCRIPTION

SCOTCH is a software package and libraries for sequential and parallel graph partitioning, static mapping and clustering; sequential mesh and hypergraph partitioning; and sequential and parallel sparse matrix block ordering.

Its main use is to subdivide a scientific problem, expressed as a graph, into a set of subproblems as independent as possible from each other (in terms of connecting edges).

FUNCTIONAL DESCRIPTION

SCOTCH takes the form of a set of libraries, plus additional standalone programs. The sequential and parallel libraries provide a set of interfaces to describe centralized and distributed graphs to partition, the target architectures to map onto, the resulting centralized and distributed mapping and ordering structures, etc. SCOTCH takes advantage of Posix threads, and its parallel version, PT-SCOTCH, uses the MPI interface.

- Participants: François Pellegrini, Cédric Lachat, Rémi Barat and Cédric Chevalier
- Partners: CNRS - IPB - Region Aquitaine
- Contact: François Pellegrini
- URL: <http://www.labri.fr/~pelegrin/scotch/>

6.5. TreeMatch

KEYWORDS: Intensive parallel computing - High-Performance Computing - Hierarchical architecture - Placement

SCIENTIFIC DESCRIPTION

TreeMatch provides a permutation of the processes to the processors/cores in order to minimize the communication cost of the application.

Important features are : the number of processors can be greater than the number of applications processes , it assumes that the topology is a tree and does not require valuation of the topology (e.g. communication speeds) , it implements different placement algorithms that are switched according to the input size.

Some core algorithms are parallel to speed-up the execution.

TreeMatch is integrated into various software such as the Charm++ programming environment as well as in both major open-source MPI implementations: Open MPI and MPICH2.

FUNCTIONAL DESCRIPTION

TreeMatch is a library for performing process placement based on the topology of the machine and the communication pattern of the application.

- Participants: Emmanuel Jeannot, François Tessier, Adele Villiermet, Guillaume Mercier and Pierre Celor
- Partners: CNRS - IPB - Université de Bordeaux
- Contact: Emmanuel Jeannot
- URL: <http://treematch.gforge.inria.fr/>

6.6. hwloc

Hardware Locality

KEYWORDS: HPC - Topology - Open MPI - Affinities - GPU - Multicore - NUMA - Locality

FUNCTIONAL DESCRIPTION

Hardware Locality (HWLOC) is a library and set of tools aiming at discovering and exposing the topology of machines, including processors, cores, threads, shared caches, NUMA memory nodes and I/O devices. It builds a widely-portable abstraction of these resources and exposes it to applications so as to help them adapt their behavior to the hardware characteristics. They may consult the hierarchy of resources, their attributes, and bind task or memory on them.

HWLOC targets many types of high-performance computing applications, from thread scheduling to placement of MPI processes. Most existing MPI implementations, several resource managers and task schedulers, and multiple other parallel libraries already use HWLOC.

- Participants: Brice Goglin and Samuel Thibault
- Partners: AMD - Intel - Open MPI consortium
- Contact: Brice Goglin
- URL: <http://www.open-mpi.org/projects/hwloc/>