

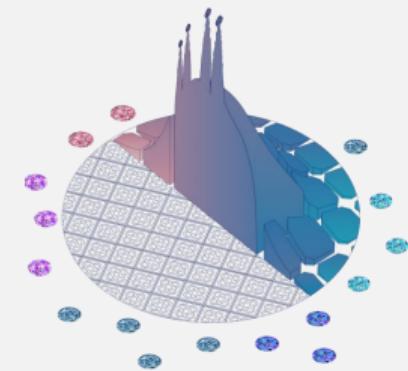
TNL-SPH: Open-source modular SPH solver for modern distributed computing platforms based on GPU accelerators

Tomáš HALADA

L. Beneš, J. Klinkovský, T. Oberhuber

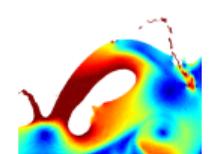
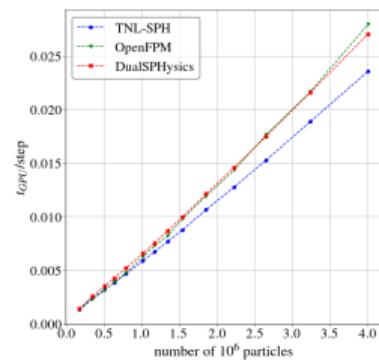
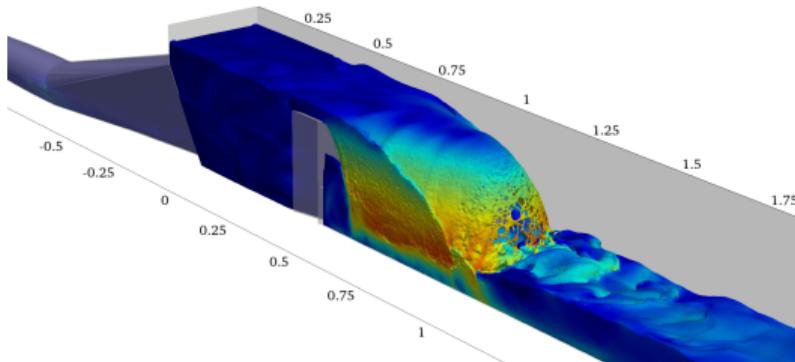
19th SPHERIC World Conference

17-19 June 2025





- We needed **multi-GPUs and distributed computing** for large scale 3D simulations in complex geometries
- **Particles are general computational algorithms** - we want general particle framework not only for SPH (*high order particle methods - LABFM, DC-PSE, GFDM, LDD, ... image representation, optimization, MD, ...*)
- Aim for better abstraction, modularity, generality and clean codebase, while trying to tip the performance



TEMPLATE
NUMERICAL
LIBRARY



LAYER 1: TNL: Architecture dependent primitives

Memory management, parallel algorithms,...

LAYER 2: Particles structure

Implementation of particle coordinates, neighbor search, particle distributed communication

LAYER 3: Particles solver

Control of the simulations, connection of particles and model variables

LAYER 4: Numerical schemes

Definition of particle fields and interaction functions



Template Numerical Library (TNL) - unified code for different hardware platforms, written in C++20, header only

- Usage of lambda functions and parallel algorithms:

```
1 const float h = 0.5;
2
3 auto v_view = v.getView();
4 auto func = [ = ] __cuda_callable__( int i ) mutable
5 {
6     v_view[ i ] = i + h;
7 };
8 TNL::Algorithms::parallelFor< Device >( 0, v.getSize(), func );
```

Unified memory management, flexible parallel reduction, parallel loops, sort,...

- Expression templates:

$$\vec{x} = \vec{a} + 2\vec{b} + 3\vec{c}$$

```
1 x = a + 2 * b + 3 * c;
```



A **general particle method algorithm** is 7-tuple

$$(P, G, u, f, i, e, \dot{e})$$

consisting of two data structures

$$P := A_1 \times A_2 \times \dots \times A_n \quad \text{the particle space,}$$

$$G := B_1 \times B_2 \times \dots \times B_m \quad \text{the global variable space}$$

such that $[G \times P^*]$ is the state space of the particle methods
and five functions

$$u := [G \times P^*] \times \mathbb{N} \rightarrow \mathbb{N}^* \quad \text{the neighborhood function,}$$

$$f := G \rightarrow \{\top, \perp\} \quad \text{the stopping condition,}$$

$$i := G \times P \times P \rightarrow P \times P \quad \text{the interact function,}$$

$$e := G \times P \rightarrow G \times P^* \quad \text{the evolve function,}$$

$$\dot{e} := G \rightarrow G \quad \text{the evolve function of global variables}$$

(J. Pahlke & I. Sbalzarini 2023)

What we understand by particle method:

$$\frac{d\mathbf{x}_i}{dt} = \sum_{j \in \mathcal{N}_i} \mathbf{f}_{\mathbf{x}}(\mathbf{x}_i, \mathbf{x}_j, \boldsymbol{\omega}_i, \boldsymbol{\omega}_j),$$

$$\frac{d\boldsymbol{\omega}_i}{dt} = \sum_{j \in \mathcal{N}_i} \mathbf{f}_{\boldsymbol{\omega}}(\mathbf{x}_i, \mathbf{x}_j, \boldsymbol{\omega}_i, \boldsymbol{\omega}_j)$$



- **Particles structure:** Structure that handles the points and neighbor search:

```
1 particles.searchForNeighbors();
```

providing the loops over the particles and neighbors:

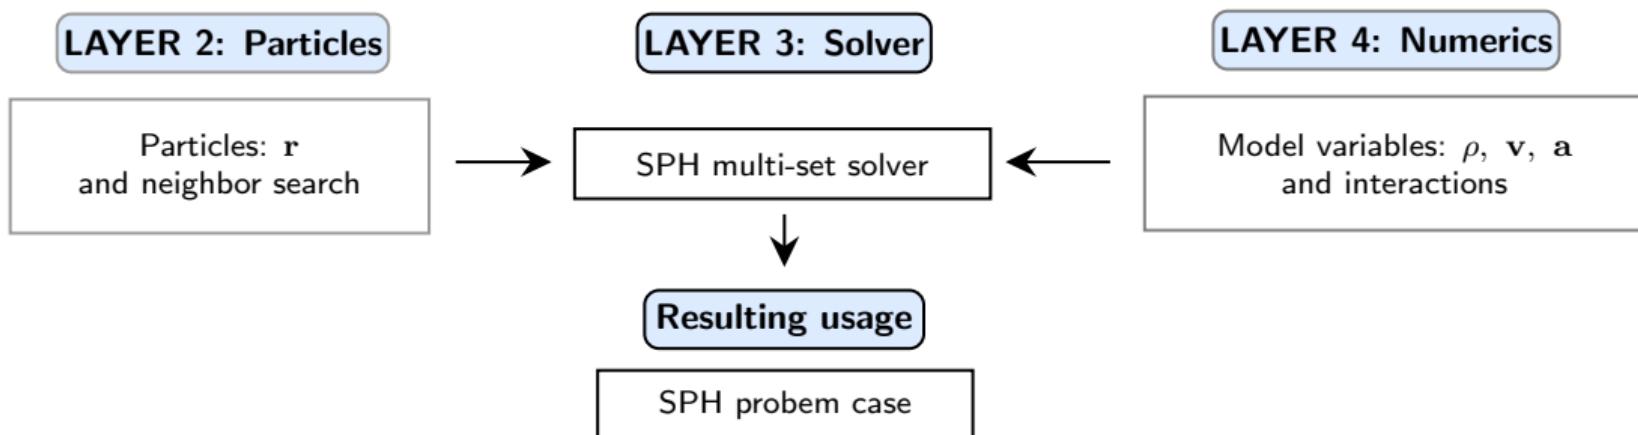
```
1 // Define operation preformed for every neighbor
2 auto pairInteraction = [=] __cuda_callable__ ( int i, int j, args... ) mutable
3 {
4 };
5
6 // Define operation performed for every particles
7 auto particleLoop = [=] __cuda_callable__ ( LocalIndexType i ) mutable
8 {
9     ParticlesType::NeighborsLoop::exec( i, r_i, searchToken, pairInteraction, args... );
10 }
11 particles.forAll( particleLoop );
```

runs in parallel even for distributed computation!

- **Particles zone:** Provides list of indices of arbitrary selected spatial subdomain.



- Multi-set solver uses **separate particle sets** for fluid, boundary particles and open boundary patches
- **Possibility to build custom time main loop** enhanced with user defined functions





Numerical solvers: Defines **variables** of the scheme/model and **interaction functions** (*fluid-fluid, fluid-boundary, boundary-fluid, pre/post interaction functions*)

```
1 auto fluidFluid = [=] __cuda_callable__ ( LocalIndexType i, LocalIndexType j, /* args i, j, ... */ ) mutable
2 {
3     // load data of neighboring particle
4     const VectorType r_j = r_view[ j ]; const VectorType v_j = v_view[ j ]; const RealType rho_j = rho_view[ j ]
5     const RealType p_j = EOS::DensityToPressure( rho_j, eosParams );
6
7     const VectorType r_ij = r_i - r_j;
8     const RealType drs = 12Norm( r_ij );
9     const VectorType v_ij = v_i - v_j;
10    const RealType F = KernelFunction::F( drs, h );
11    const VectorType gradW = r_ij * F;
12
13    // continuity equation
14    const RealType psi = DiffusiveTerm::Psi( rho_i, rho_j, r_ij, drs, diffusiveTermsParams );
15    const RealType diffTerm = psi * ( r_ij, gradW ) * m / rho_j;
16    *drho_i += ( v_ij, gradW ) * m - diffTerm;
17
18    // momentum equation
19    const RealType p_term = ( p_i + p_j ) / ( rho_i * rho_j );
20    const RealType visco = ViscousTerm::Pi( rho_i, rho_j, drs, ( r_ij, v_ij ), viscousTermParams );
21    *a_i += ( -1.0f ) * ( p_term + visco ) * gradW * m;
22}
```

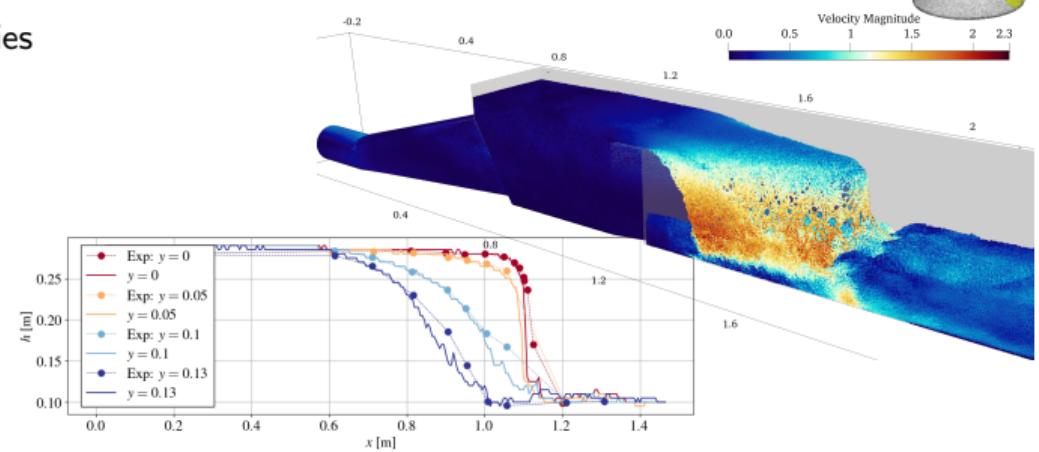
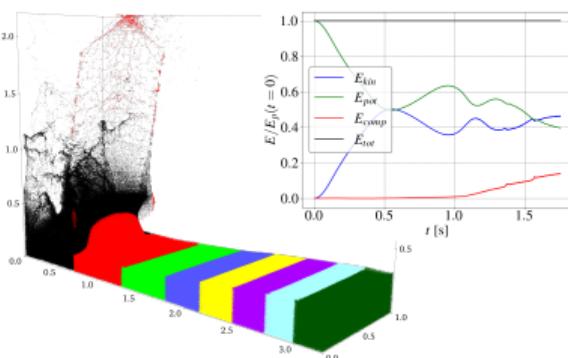
$$\frac{D\rho_i}{Dt} = \rho_i \sum_{j=1}^N (\mathbf{u}_i - \mathbf{u}_j) \cdot \nabla_i W_{ij} V_j + \delta_\psi h c_0 \mathcal{D}$$

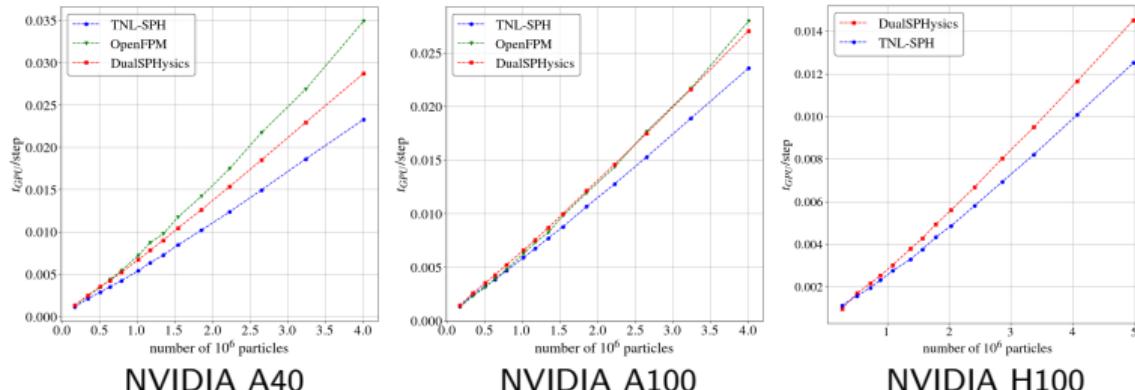
$$\frac{Du_i}{Dt} = -\frac{1}{\rho_i} \sum_{j=1}^N (p_i + p_j) \nabla_i W_{ij} V_j + \frac{1}{\rho_i} \sum_{j=1}^N \Pi_{ij} \nabla_i W_{ij} V_j + \mathbf{f}_i$$



Implementing particular schemes should not be considered as a problem,...

- WCSPH (diffusive terms, various boundary conditions: DBC, MDBC, **boundary integrals**)
- Unconditionally stable SPH scheme (*Cercos-Pita, 2024*)
- RSPH, **SHTC-SPH** - unified framework for solid and fluids
- various integration schemes (Verlet, symplectic Verlet, RK45, implicit midpoint)
- open boundary conditions (inlet, outlets)
- discretization of external geometries

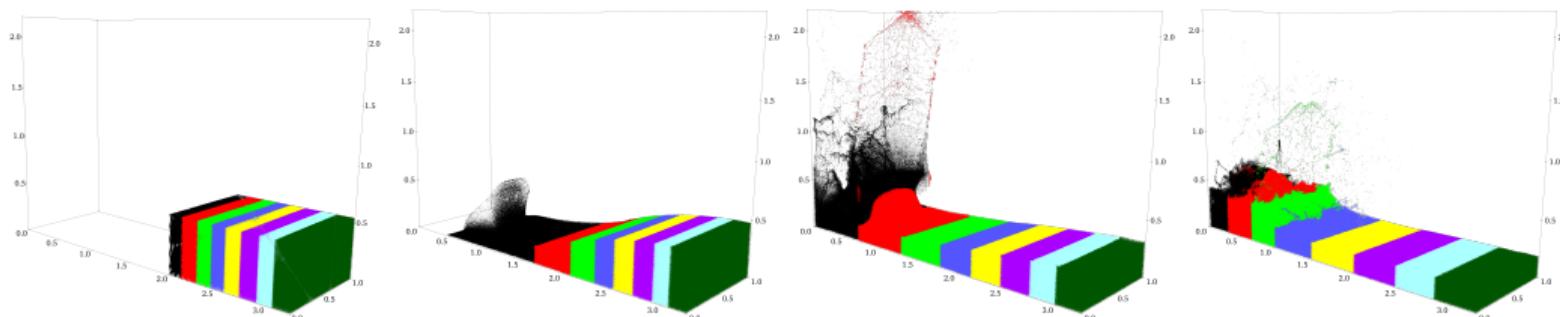
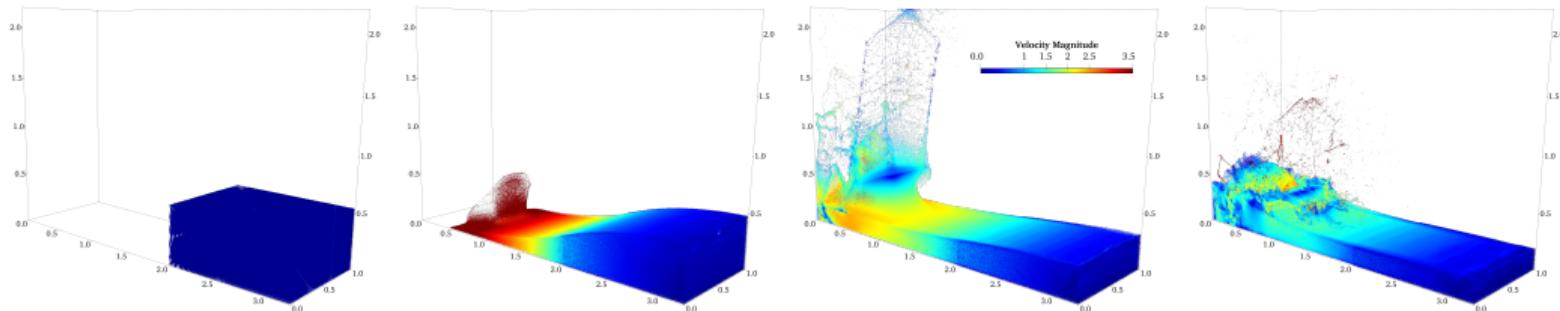




GPU	TNL-SPH		DualSPHysics		OpenFPM	
	t_{step} [s]	MPUPS	t_{step} [s]	MPUPS	t_{step} [s]	MPUPS
A40	0,023	172,3	0,029	139,6	0,035	114,9
A100	0,024	170,2	0,027	148,2	0,028	143,4
H100	0,013	320,5	0,015	276,2		

Test case: SPHERIC testcase 2: 3D dam break, δ -WCSPH scheme with Molteni-DT, artificial viscosity, DBC, Verlet scheme, $\Delta t = const.$, resolutions $\Delta x = [0.02, 0.006]$, **table shows data for $\Delta x = 0.006$**

Benchmark setups for all solvers: <https://github.com/tomashalada/sph-benchmarks/>



Support of multi-GPUs and distributed computations using MPI, with 1D, 2D or 3D domain decomposition and subdomains load balancing.



Strong scaling:

N_{GPUs}	GPUPS	$t_{\text{step}} [\text{s}]$	Speedup	Eff
1	0,297	0,346	1,00	1,00
2	0,585	0,176	1,97	0,99
4	1,157	0,089	3,90	0,97
8	2,154	0,048	7,26	0,91

Weak scaling:

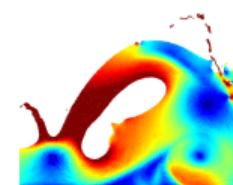
N_{GPUs}	$N_{\text{ptcs}}(10^6)$	GPUPS	$t_{\text{step}} [\text{s}]$	Speedup	Eff
1	15,2	0,335	0,045	1,00	1,00
2	28,5	0,622	0,046	1,98	0,99
4	53,8	1,147	0,047	3,87	0,97
8	102,8	2,136	0,048	7,54	0,94

Test case: SPHERIC test case 2: 3D dam break, δ -WCSPH scheme with Molteni-DT, artificial viscosity, DBC, Verlet scheme, **weak scaling** 20×10^6 particles per GPU, **strong scaling** 10^8 particles, NVIDIA H100 cards



- *Iterating by particles*: The limiting factor is the memory access - too many neighborhood particles with too much data to read
- *Iterating by pairs*: The limiting factor is writing to memory - too much data to write

TNL-SPH: Open source, modular, header only, and pretty fast SPH solver oriented towards industrial hydrodynamics



TEMPLATE
NUMERICAL
LIBRARY

This research was supported by Ministry of Education, Youth and Sports of the Czech Republic under the project CZ.02.01.01/00/23021/0008954 "Hydrodynamic machines for smart energetics"



Co-funded by
the European Union



MINISTRY OF EDUCATION,
YOUTH AND SPORTS