Towards More Efficient Implementations of Multiscale Thermal-Hydraulics

Antoine Gerschenfeld

Version of Record: https://www.sciencedirect.com/science/article/pii/S0029549321002740
Manuscript_0b6dc0e735b79f4695325ea4bb7c87a6

Towards More Efficient Implementations of Multiscale Thermal-Hydraulics

Antoine Gerschenfeld

Université Paris-Saclay, CEA, Service de Thermo-hydraulique et de Mécanique des Fluides, 91191, Gif-sur-Yvette, France

Abstract

Over the past decade, the need to supplement system-scale simulations of reactor transients with the results of finer simulations (subchannel or CFD) has increased continuously. In many cases, the local phenomena predicted at these scales (such as flow patterns in the core or within inlet/outlet plenums) can affect the overall transient: in that case, then all codes should be run concurrently in a consistent manner in order to obtain a single, "multi-scale" simulation of the transient of interest.

Because their subchannel/CFD components tend to require meshes beyond the capabilities of the 3D modules present in modern system codes, most multiscale simulations can only be performed by coupling different codes together. The strategy used to implement this coupling can have a crucial impact on both the solution accuracy and on the numerical cost of the calculation : in particular, algorithms which require small time steps or large number of iterations between the codes can multiply the numerical cost of multiscale compared to an (already expensive) standalone CFD simulation.

This paper discusses a range of algorithms suitable for coupling thermal-hydraulics codes at either thermal or hydraulic boundaries. These algorithms are grouped into four broad classes of increasing complexity (fixed-point, improved fixed-point, quasi-Newton and Newton). The more complex variants are more difficult to implement, but have been observed to significantly decrease the numerical overhead of multi-scale coupling.

Keywords: thermal-hydraulics, multiscale, code coupling, STH, CFD

1. Introduction

The safety justification of existing and prospective nuclear reactors requires, for its thermal-hydraulics component, simulation tools able to represent all phenomena and components of interest in the primary circuit, and possibly to other circuits as well (power conversion system, intermediate and decay-heat removal loops, etc.). Since the 1970s, system thermal-hydraulics (STH) codes such as CATHARE, RELAP or TRACE have been developed and validated to fulfill this need: by modelling each circuit as a network of 0D, 1D and coarse 3D elements combined with specific component models (for fuel assemblies, heat exchangers, pumps, etc.), they offer a validated and numerically efficient way to simulate even the longest transients (going into the 10^5s). As such, STH simulations have played a key role in the safety justification of almost all current reactors.

Because the simulation scale adopted in STH is relatively coarse, all local phenomena that may impact a given transient of interest must be described by physical models: however, establishing and validating these models can be difficult. For example:

- In pressurized-water reactors (PWRs), the reactor's behavior during dissymmetric transients, such as loss-of-coolant accidents (LOCA), steam generator tube ruptures (SGTR), or steam line breaks (SLB) will be strongly affected by how flow propagates from the cold legs to the core inlet (via the downcomer and inlet plenum) and from the core outlet to the hot legs (via the outlet plenum) (fig. 1). In forced convection, these mixing effects can be modelled in STH via experimentally-validated "mixing matrices"; however, at lower flow rates (or if the flow reverses), flow patterns within these regions tend to become so complex that the mixing matrix approach is no longer reliable[1].

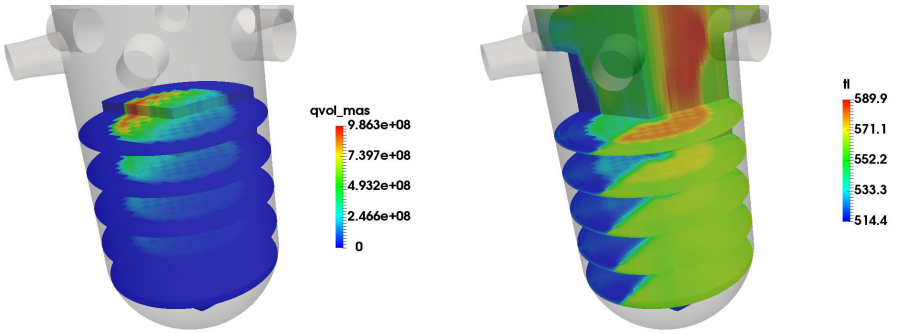

Figure 1: Simulated power (left) and temperature distribution (right) in an APR1400 PWR during a SLB transient[1] . Overcooling in the affected SG leads to dissymmetric inlet temperatures in the cold legs: this dissymmetry propagates into the core, where it leads to dissymmetric power and temperature profiles.

- Sodium fast reactors (LMRs) often adopt a pool-type primary circuit design in order to minimize LOCA risks and increase thermal inertia. In these designs, decay heat removal often relies on passive natural convection in the primary circuit: however, in such a regime, the large sodium plena around the core stratify, while the jets at the outlet of the core and heat exchangers transition from velocity- to buoyancy-driven flow (fig. 2). These phenomena have such a strong impact on natural convection that neglecting them can lead to overestimating the core flowrate by as much as 50%[2]. Additionally, 3D convection loops between the core and its outlet plenum can contribute to up to 30% of the overall decay heat removal.

The traditional way to account for such effects at the system scale would be to introduce conservative hypotheses in the modelling (for instance, by postulating mixing effects in the outlet plenum to be as unfavorable as possible) and to increase design margins as needed. More recently, thanks to advances in numerical simulation, it has become possible to simulate a number of these phenomena : the examples described above can now be simulated using RANS-scale CFD (within free regions) and subchannel-scale or porous-medium CFD (within the core) at a numerical cost compatible with the long durations associated with

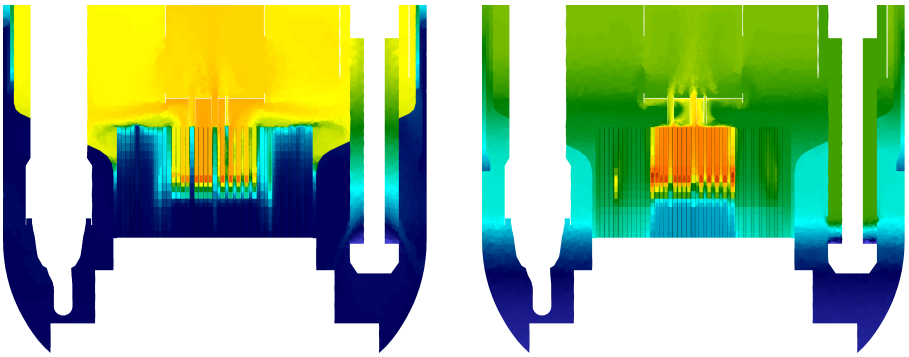

Figure 2: Impact of 3D phenomena on decay heat removal by natural convection in SFRs[2]. During the transition to natural convection, the hot and cold pool transition from their initial homogeneous states (left) to stratified regimes (right), while the outlet jets from the core and heat exchangers become buoyancy-driven. The natural-convection regime is influened by heat exchanges between the hot and cold pool as well as by lateral cooling of the core assemblies.

reactor-scale transients.

Applying these new capabilities to reactor-scale simulations requires either the development of a new code or the integration of finer scales into existing STH simulations. The second option has proven particularly attractive, as it allows one to capitalize on the efforts expended to develop and validate models at the STH scale (such as pumps, heat exchangers and point kinetics) : it also allows for a progressive approach to the safety justification, in which a system-scale model suitable for forced-convection transients can be augmented "on-demand" with finer scales when simulating more complex transients.

However, the system-scale, subchannel-scale and CFD-scale parts of such a multiscale simulation will often rely on different existing codes. Integrating them into a consistent simulation of the complete reactor requires developing a coupling algorithm. The choice of this algorithm is crucial to the accuracy and performance of the multi-scale simulation itself: errors from imprecise coupling may accumulate as the transient progresses, while constraints on the maximum possible timestep, or requirements for high number of iterations between the codes, may increase the numerical cost of the simulation by an order of magnitude compared to a standalone CFD (which is often already much more

expensive than STH).

Because they are strongly affected by 3D effects that can be simulated by single-phase CFD, SFRs and other liquid-metal cooled reactors have been the focus of strong efforts in recent years, especially within the EU projects THINS[3, 4] and SESAME [5, 6, 7]; thanks to these efforts, multi-scale simulations reached a high enough level of maturity to become the reference safety analysis for natural-convections transients in the ASTRID[8] 600-MWe French SFR project. On the LWR side, efforts in the USA[9], in Korea[1] as well as within the NURESIM project[10] have contributed to advance the state of the art : multi-scale simulations have not yet become the reference option for a LWR safety transient, but may achieve this status in the near future.

The aim of this article is to assist in the development of multi-scale simulations by presenting and discussing a range of coupling algorithms, based on the experience accumulated in their practical implementation. The first section discusses overall strategies for coupling thermal-hydraulics codes and the general requirements one may impose on coupling algorithms; then, in a second section, several such algorithms are presented and discussed. These algorithms are grouped into four broad categories (fixed-point, improved fixed-point, quasi-Newton and Newton): successive categories tend to be more difficult to implement, but bring additional benefits in terms of accuracy and/or performance. This overview is followed by an example illustrating how, in the case of SFR modelling at CEA, the need to model additional physical phenomena made it necessary to consider new coupling algorithms. Finally, a general conclusion discusses expected evolutions in the state of the art.

2. Strategies and requirements

The first step in the development of a multi-scale simulation is to identify, for a given transient, which phenomena should be simulated via fine-scale (sub-channel or CFD) simulation. This leads to the identification of one or more *fine domains* associated with these phenomena, which will become the calculation

domains of the subchannel/CFD codes. In the introductory examples of Section 1, CFD may be used to model the downcomer and inlet plenum of a LWR or the hot and cold pools of and SFR, while the subchannel scale may be used to model the complete core at either the subassembly scale (one mesh per S/A) or the subchannel scale.

Once these domains are chosen, one then must choose how to treat the parts of the initial STH model that are to be superseded by finer simulations. Two strategies are available:

- in the *domain decomposition* strategy (fig. 3.a), the sections of the STH model covered by fine domains are removed, so that each part of the overall computation is only covered by one code;
- in the *domain overlapping* strategy (fig. 3.c), the STH model is left intact: the regions of the overall domain computed by the fine scale are thus also computed at the system scale.

In order to obtain a consistent overall calculation, data exchanges must take place between the different codes. In the domain decomposition approach, these exchanges can only take place at the edges of the fine domains, where the domains of the STH and fine scale intersect; in the case of domain overlapping, data exchanges may also be needed between the "bulk" of the fine domain and its STH counterpart. The discussion of the coupling algorithms in Section 3 will lead to a discussion of the pros and cons of these two approaches in the conclusion.

In both of these approaches, two main categories of coupling boundaries can be identified:

- *Thermal boundaries*, which are only subjected to heat transfers: for instance, between a CFD model of the primary side of an exchanger and an STH model of the exchanger tubes.
- *Hydraulic boundaries*, which are crossed by fluid flow (and across which heat exchanges are usually neglected): for instance, in a PWR, between a

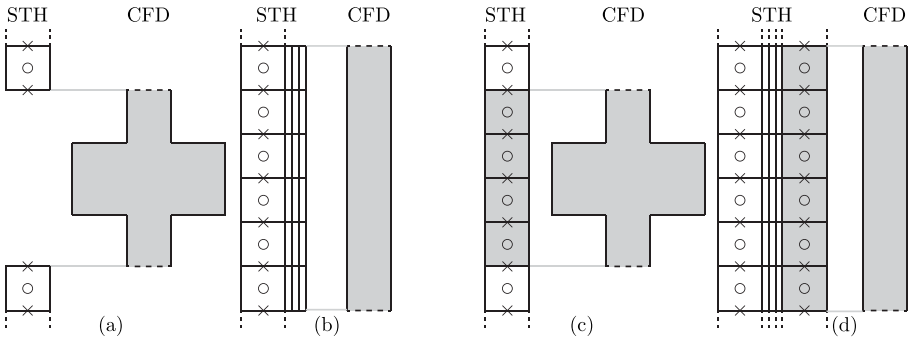

Figure 3: Examples of hydraulic and thermal coupling boundaries in the domain-decomposition (left) and domain-overlapping (right) approaches. In the case of a hydraulic boundary, the fine domain is either removed from the system scale (a) or "overlapped" by the the fine scale (c); for a thermal boundary, the STH model of an heat exchanger can either interface directly with CFD (b) or with an overlapped system region (d). We assume that the system code uses a staggered mesh scheme, where scalar unknowns void fraction, temperature, pressure) are located at cell centers (denoted by circles) while velocity unknowns are located at cell faces (crosses)

STH model of a cold leg pipe and a CFD model of the downcomer.

When designing algorithms for exchanging data at these boundaries, it can be useful to reason in terms of conservation properties. At each time step, given amounts of mass, energy and momentum will cross each side of a given coupling boundary: for the overall simulation to be valid, the amounts entering/exiting the system code must be as close as possible to those exiting/entering the fine model. In practice:

- *mass conservation* is of paramount importance: as most reactor transients consist in simulating closed circuits over large periods of time, errors in mass conservation will tend to accumulate and lead to a non-compensated "drift" of the overall simulation over time;

- *energy conservation* is of slightly smaller importance but should still be tracked closely, as errors will usually result in an artificial recurring heat sink/source at the coupling boundaries and will therefore affect the overall

temperature distribution;

- *momentum conservation* (which is equivalent to guaranteeing that the pressure field is continuous at the coupling boundaries) is less important in forced convection, but becomes crucial in natural convection, where slight errors may overcome the small driving forces generated by buoyancy.

The algorithms described in the next section aim to guarantee as much as possible each of these conservation relations in decomposition-type and overlapping-type multiscale couplings. In most cases, these relations can only be maintained by using an implicit time scheme in which, at each time step, iterations between the codes are required to converge to a common state. If the coupling algorithm requires both small time steps (for stability) and a high number of internal iterations, then the "coupling overhead" (the time taken by the coupled simulation compared to running the models separately) will be particularly high.

3. Coupling algorithms

The algorithms presented in this sections have been grouped into four broad categories:

- *fixed-point* algorithms, where parameters of interest at coupling boundaries (such as flow rate, pressure, temperature and heat fluxes) are exchanged between the codes. Using these exchanges, the codes are iterated in a simple fixed-point loop until they converge to a common value of these parameters for a given time step;

- *improved fixed-point* algorithms, where, on top of the parameters of interest, additional values are exchanged between the codes in order to obtain faster convergence;

- *quasi-Newton* algorithms, in which the fixed-point iteration of the first two categories is replaced by a higher-order algorithm (typically Newton-Raphson), while relying on the same data exchanges;

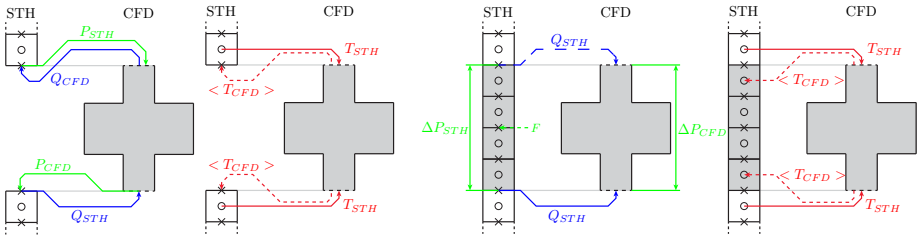

Figure 4: Fixed-point algorithms for coupling hydraulic boundaries in the domain-decomposition (left) and domain-overlapping (right) approaches.

- *monolithic Newton* algorithms, where even more data is exchanged between the codes to ensure that the code-to-code iterations follow (or are equivalent to) a Newton algorithm.

3.1. Fixed-point algorithms

These algorithms are generally the first step in the implementation of a multi-scale coupling. Because they only rely on exchanging "standard" input and output variables (flowrates, pressures and temperatures), they generally require few developments in the codes to be coupled together, namely:

- an interface to pilot the calculation (choice of time step, execution of iterations, time advancement);
- an interface to set input quantities and obtain output variables in-between these steps.

Provided these capabilities are available, coupling algorithms below can be implemented.

Hydraulic boundaries

At hydraulic boundaries (fig. 4), mass, momentum and energy conservation must be implemented. In the domain-decomposition approach (left), mass and momentum can be synchronized by imposing a pressure boundary condition on one side of the boundary and an imposed flowrate boundary condition on

the other side. From a coupling perspective, both choices are equally suitable regardless of the flow direction. At each iteration, the calculated result on one side (flowrate for a pressure-type boundary condition and vice versa) is then used as a boundary condition on the other side: this process is iterated until both codes converge to common values.

It should be noted that, if an imposed-flow BC is used on the CFD side, then a velocity profile must be chosen to represent the average velocity provided by the STH side. The naive choice of a uniform profile will usually lead to an artificial "flow development effect" in the nieghborhood of the coupling boundary: this can be avoided by imposing a prescribed profile approximating a fully-developed flow. Alternatively, if an imposed-pressure boundary condition is chosen on the CFD side, then no velocity profile needs to be prescribed.

Reactor simulations often include large single-phase, quasi-incompressible portions, which induce strong long-range coupling in the pressure field: in that case, the pressure/flowrate fixed-point iterations described above tend to become unstable at large timesteps and/or require tens to hundreds of iterations.

For quasi-incompressible flow, this pain point can be avoided by using the equivalent algorithm for domain-overlapping coupling (fig. 4, right): in that case, mass conservation can be ensured by imposing the STH-computed flowrate at $N-1$ of the N boundaries of the CFD domain, and flowrate equality at the last boundary will be obtained as a consequence of incompressibility in both codes. To ensure momentum conservation, it is then sufficient to guarantee that, at the end of each timestep, the pressure *difference* between any two hydraulic boundaries in the STH simulation will follow the value obtained on the CFD side : this can be ensured by adding momentum source terms on the inside of the STH side and adjusting them until ΔP_{STH} converges to ΔP_{CFD}. Because this process iterates over pressure differences rather than over the pressure field itself, it does not suffer from the instability and high iteration counts that often plague the decomposition approach; on the other hand, its restriction to incompressible flows means that it is clearly unsuitable to two-phase simulations.

Finally, in both cases, energy conservation must be ensured by adjusting

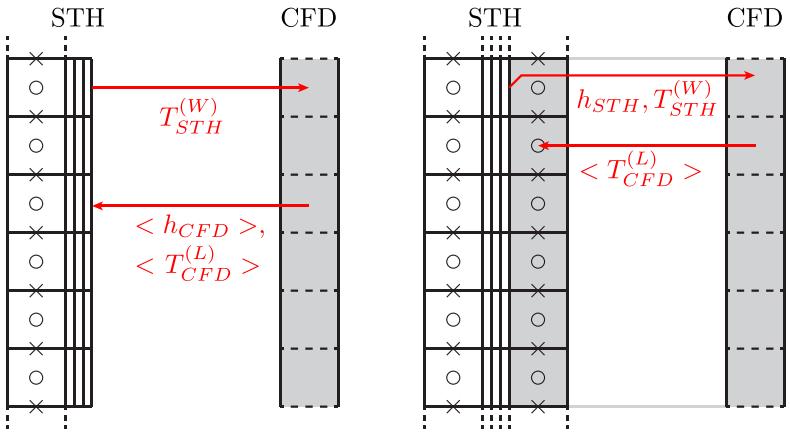

Figure 5: Fixed-point algorithms for coupling thermal boundaries in the domain-decomposition and domain-overlapping approaches.

inlet/outlet temperatures at both sides of the boundary. On the CFD side, the single temperature value computed on the STH side can be imposed as a boundary condition; on the STH side, the value to be imposed is a velocity-weighted average of the boundary temperature profile computed by the CFD. This value must be imposed either as a STH boundary condition in the decomposition approach or on the meshes right inside the overlapped domain in the overlapping approach: then, provided that the STH code uses an upwind convection scheme in its energy equations, the energy exiting the CFD domain will be consistent with that entering the STH domain.

Thermal boundaries

At thermal boundaries (fig. 5), such as the primary side of a heat exchanger, the STH-computed wall-temperature $T^{(W)}_{STH}$ can be interpolated onto the corresponding CFD meshes and used as input to a volumetric heat source term of the form

$$\Phi_{CFD} = h(T^{L}_{CFD} - T^{W}_{STH}),$$

with T^{L}_{CFD} the (local) CFD fluid temperature in each mesh. The heat transfer coefficient h can be computed by the CFD code (usually via a correlation for

Nusselt number); in the domain-overlapping approach (fig. 5, right), it can also be taken from the STH code. On the STH side, a boundary condition of the form

$$\Phi_{STH} = h(< T^L_{CFD} > -T^W_{STH})$$

must then be implemented in the domain-decomposition approach, with $< T^L_{CFD} >$ the average of the fluid CFD temperature over each STH mesh; in the domain-overlapping case, this average temperature can instead be imposed directly inside the (overlapped) STH fluid domain. Iterations then must be performed to allow $< \Phi_{CFD} >$, the average CFD heat flux, to converge to a common value with Φ_{STH}.

Compared to alternative approaches involving direct heat flux exchanges, this temperature-based fixed point avoids two potential pitfalls:

- instability at high time steps, which can make the fixed-point iterations unstable;
- unphysical temperatures in the CFD domain: when exchanging fluxes, if the fluid flow is slow or stagnant in part of the CFD, then imposing a constant, average heat flux may lead to some local temperatures T_{CFD} falling below T_{STH}.

However, obtaining heat flux convergence can require high iteration counts, especially if the thermal conductivity is very high (such as in sodium).

3.2. Improved fixed-point algorithms

This category groups two algorithms which, while using the same type of fixed-point iterations used in the previous section, rely on transmitting more parameters between the codes in order to obtain improved convergence. It includes:

1. A *quasi-implicit* coupling algorithm at thermal boundaries (fig. 6, left). This algorithm relies on extracting from the STH code both the heat flux Φ_{STH} and its sensitivity to the fluid temperature $\partial\Phi_{STH}/\partial T^{(L)}_{STH}$ (in

the domain-overlapping approach) or to the external temperature boundary condition $\partial\Phi_{STH}/\partial T_{EXT}$ (in the domain-decomposition approach). These coefficients are usually computed internally by STH codes as part of their internal solution methods: if they can be passed to the CFD, then they can be used to compute the heat flux Φ_{CFD} and its average over each STH mesh $< \Phi_{CFD} >$ without stability issues. Then, at the next STH iteration, the STH heat flux can be set directly to $< \Phi_{CFD} >$ in order to ensure energy conservation : hence, two STH iterations and a single CFD iteration are sufficient to obtain a converged heat flux at the boundary.

2. A similar algorithm to couple, in the overlapped approach, a STH pump model in the overlapped domain to the CFD. In simulations where the CFD domain covers a complete circuit, such as the primary pool of an SFR, a model for the pressure head provided by the circuit pumps must be implemented. STH codes usually employ "point" pump models, where the hydraulic head of the pump, H_{pump}, is expressed as a function of the shaft rotation speed ω_{pump} and the pump volumetric flowrate Q_{pump}:

$$H_{pump} = f(\omega_{pump}, Q_{pump}) \,,$$

together with a point equation on ω_{pump} for the angular momentum balance on the pump shaft.

Instead of reimplementing such a model from scratch on the CFD side, one may reuse the STH pump model: the pump head H_{STH} is extracted from the system code and imposed as a momentum source term on the CFD side, while the flowrate computed across the CFD surface corresponding to the pump, Q_{CFD}, is imposed on the STH side at the next iteration.

In practice, this "naive" coupling scheme turns out to be unstable at most time steps, especially in cases where several pumps are present in parallel (the common case in most primary circuits). However, stability can be improved by obtaining from STH the sensitivity coefficient $\partial H_{STH}/\partial Q_{STH}$: it can then be implemented on CFD side as a velocity-dependent source term. This sensitivity coefficient is inherently stabilizing (higher velocities

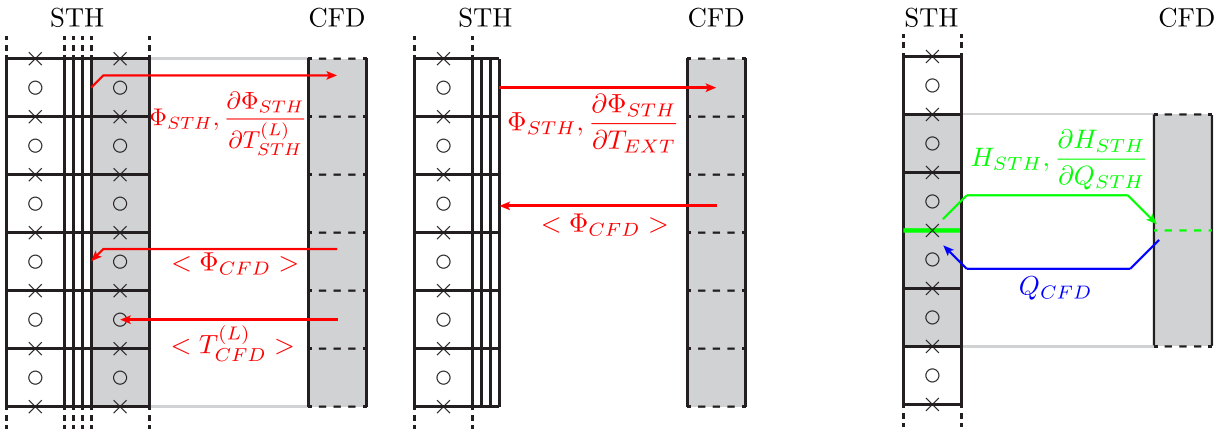

Figure 6: Improved fixed-point algorithms for thermal boundaries in the domain-overlapping and domain-decomposition approaches (left) and for coupling a pump in the domain-overlapping approach (right).

will lead to a lower source term): in practice, convergence can usually be observed in two or three iterations.

3.3. Quasi-Newton algorithms

The algorithms used in the two sections above can be used to implement an efficient multi-scale coupling (few iterations, no time-step restrictions):

- at thermal boundaries in both the domain-decomposition and the domain-overlapping cases;
- at hydraulic boundaries in the domain-overlapping case, provided the flow is quasi-incompressible.

This leaves the case of a hydraulic coupling in the domain-decomposition formulation. In order to overcome the low stability/high iteration counts associated with the pressure/velocity coupling of fig. 4 (left), a potential solution is to replace the standard Gauss-Seidel fixed-point iteration procedure with a faster-converging method. Aside from accelerated fixed-point methods such as Aitken's, Newton's method is an especially attractive candidate: however, it

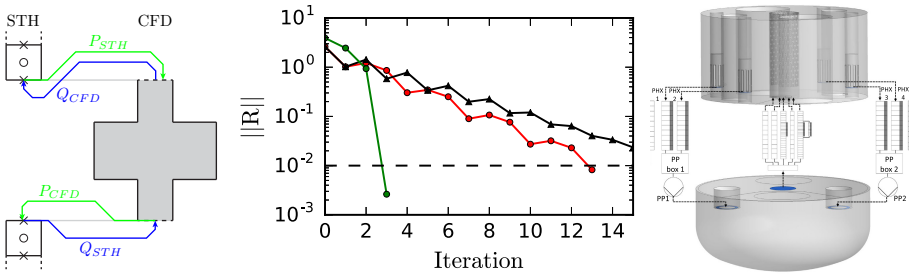

Figure 7: Improvement of a domain-decomposition hydraulic boundary coupling through a quasi-Newton iteration algorithm: simplified configuration (left), pressure convergence (center) for the quasi-Newton approach (green) compared to Gauss-Seidel fixed-point (black) and Aitken fixed-point (red)[11], configuration for a multi-scale simulation of the MYRRHA reactor[12] (right).

requires the construction of a Jacobian matrix of sensitivity coefficients at each iteration.

In [11, 12], this method was implemented in a STH/CFD coupling between RELAP5-3D and FLUENT without extracting additional information from the codes (fig. 7). In order to compute the Jacobian matrix, discrete derivatives are computed over the response function of each code: for the CFD code in the configuration of fig. 4, this function takes the form

$$(Q_{CFD}^{(top)}, P_{CFD}^{(bottom)}) = f_{CFD}(Q_{STH}^{(bottom)}, P_{STH}^{(top)})$$

Discrete derivatives can be used to compute the Jacobian of f_{CFD} by evaluating the CFD code five times; once this is done as well for the STH code, convergence can be accelerated to just a few iterations (fig. 7, center), compared to more than 10 iterations for fixed-point approaches.

While this technique provides greatly accelerated convergence, its main weak point lies in the number of iterations required to obtain the Jacobian matrix. This process imposes a certain minimum number of iterations, which moreover grows quadratically with the number of hydraulic pressure boundaries. Two techniques can be implemented to mitigate this cost:

- instead of updating the Jacobian at each time step, one may attempt to

reuse its value at the previous timestep[11]. This introduces an approximation in the solution algorithm: if the old Jacobian is no longer up to date, then convergence may not be accelerated. This leads to an on-demand approach where the Jacobian is reused as much as possible, then only updated when slow convergence is observed;

- in complex problems such as (fig. 7, right), where the CFD domain may contain as many as 10 outlets, the size of the Jacobian can be reduced by only considering groups of input/output variables. For instance, the model presented in [12] includes five hydraulic boundaries representing outlets of subassembly groups: in that case, only the "total flowrate" variable $Q_{core} = \sum_1^5 Q_{core,i}$ was considered when building the Jacobian. In that manner, the size of the Jacobian matrix was reduced from 10x10 to 4x4, leading to a 6-fold reduction in the number of discrete derivatives. The quality of this approximation was checked by confirming that this approximate Jacobian still led to significantly accelerated convergence.

3.4. Monolithic Newton algorithms

The results of the quasi-Newton approach (Sec. 3.3) highlight the potential of Newton iterations to accelerate convergence in multi-scale simulations, provided that the Jacobian matrix can be assembled efficiently. Instead of using discrete derivatives, one may choose to modify the codes so that they provide, in addition to their output variables, the sensitivity coefficients of these output variables relative to the code's input boundary conditions. If that is possible, then two types of Newton algorithms can be envisioned:

1. a *symmetric* algorithm, in which both codes provide similar information so that the coupling algorithm itself can run the Newton algorithm. In that configuration :
   (a) all codes use imposed-pressure boundary conditions at the coupling boundaries;

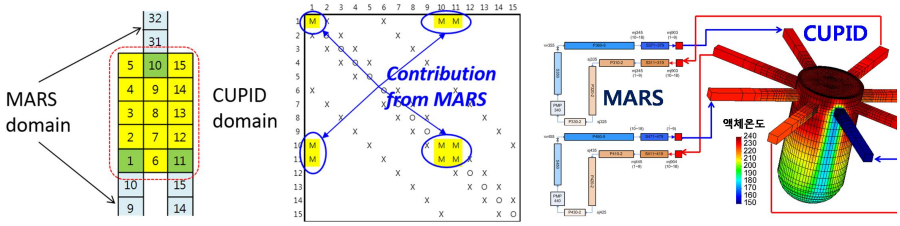

Figure 8: Monolithic Newton coupling between the MARS STH code and the CUPID CFD code at KAERI[1]. At each iteration, the sensitivities of the MARS output parameters are integrated into the CUPID solution matrix: then, once the CUPID iteration is complete, MARS is updated using the new boundary variables. This coupling was applied[13] to a ROCOM flow-mixing test[14].

(b) at each iteration, each code not only provides the flowrate Q_i going through the boundary i, but also its derivatives with respect to the pressures at the other boundaries $\partial Q_i/\partial P_j$;

(c) using this information, the coupling algorithm assembles a Jacobian matrix for the continuity equations $Q_i^{(in)} = Q_i^{(out)}$ and solves for the boundary pressures P_j;

(d) these pressures are substituted back into each code before the next iteration.

2. alternatively, a master/slave approach can be used, in which one code first computes its sensitivity parameters: those are integrated as boundary conditions into the solution algorithm of the other code, which can then compute an iteration on its own. Finally, the values obtained for the boundary variables are used by the first code to update its internal variables.

Both methods will give identical results: additionally, these results will be the same as if the resolution of both codes was performed using a single linear system containing both codes' internal variables. As such, these schemes can be categorized as "monolithic Newton" : the processes described above amount to a Schur-complement solution of the overall matrix. The symmetric method (1) is in fact rather similar to the solution method used within most STH codes:

for instance, in the CATHARE code, at each iteration,

1. every 0D/1D/3D element provides output flowrates at each of its junctions with the rest of the network, as well as their sensitivities, to a circuit-level Jacobian matrix;
2. this matrix is then solved in order to update the junction variables (mainly pressures and temperatures);
3. then, all elements update their internal variables using the new junction variables.

The master/worker form (2) of this monolithic coupling has been implemented by KAERI[13, 1] between the MARS system code and the CUPID CFD code (fig. 8). In this coupling, the MARS code provides additional contributions to the CUPID pressure matrix, which can then be solved normally: afterwards, the MARS code updates its internal variables from the new boundary pressures. A similar coupling is under development at CEA between the CATHARE system code and a hybrid domain combining the TrioMC subchannel code and the TrioCFD CFD code: the current plan is to integrate TrioMC/TrioCFD into the CATHARE element/junction solution procedure, which would result in a master/worker architecture (2).

It should be mentioned that, in some cases, it may prove more straightforward to assemble a full Jacobian matrix integrating the internal variables of every code, and then solve it in a single pass. This procedure has proved to be the most efficient when coupling solid thermal conduction to thermal-hydraulics within the TrioCFD code.

4. Example: Multi-scale modelling of a small-scale SFR

In order to illustrate the tradeoffs between different coupling algorithms, this section presents the multiscale modelling at CEA of a design for a reduced-power (400 MWth) pool-type SFR (compared to the 1500 MWth ASTRID design shown on figure 2). The coupling approach developed for SFRs at CEA, validated on the PHENIX reactor [3] and applied to ASTRID, consists in:

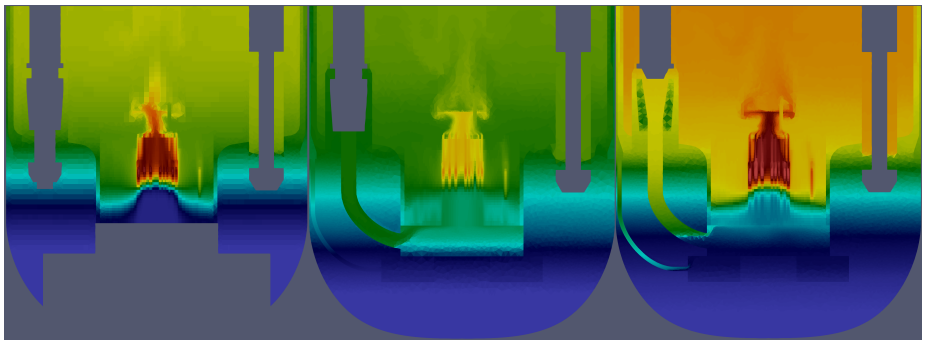

Figure 9: CEA multiscale model of a 400 MWth pool-type SFR : initial model with hot/cold pool CFD (left), intermediate model with CFD from pump outlet to core inlet (center), final model with complete primary pool CFD (right).

- a system-scale CATHARE model of all circuits (primary, intermediate and decay heat removal);
- a TrioCFD CFD model of a domain encompassing the inter-wrapper region, the hot and cold pools, and the primary sides of the IHXes;
- a TrioMC subchannel model of the core.

The TrioMC and TrioCFD models overlap the CATHARE primary circuit model using the fixed-point overlapping algorithm (Sec. 3.1); in each IHX, the CFD primary side is coupled to the CATHARE secondary side using the thermal improved fixed-point algorithm (Sec. 3.2). These choices lead to low iteration counts, while avoiding the use of CFD in zones of the primary circuit without significant 3D effects : the primary pumps, the diagrid (core inlet plenum) and the pump-diagrid pipes connecting them are not included in the CFD domain, and so are only modelled at the system scale.

The initial model for the reduced-power, 400 MWth SFR (fig. 9, left) followed this approach. However, it soon became apparent that, contrary to ASTRID, this new design featured longer pump-diagrid pipes so as to shorten the primary pumps: in turn, this meant that heat exchanges between these pipes and the surrounding cold pool, which had been unimportant in ASTRID,

could no longer be neglected.

In order to account for these exchanges, a new model extending the CFD to the diagrid and pump-diagrid lines was investigated (fig. 9, center). This model used the same coupling scheme as the initial coupling, but the CFD extension allowed for the modelling of heat transfers across the (CFD) surfaces of the pump-diagrid lines, thus fulfilling the original goal of this new model.

However, in this model, the "non-overlapped" part of the STH domain is reduced to a small region around each pump: this turns out to be detrimental to the robustness and stability of the fixed-point overlapping coupling scheme (Sec. 3.1), as small errors in the convergence of ΔP_{STH} to ΔP_{CFD} in the large overlapped domain can lead to large relative ΔP errors in the much smaller STH-only domain. This in turn can lead to large errors on the flowrate: this is especially the case in natural convection, where pressure convergence errors as small as 10 Pa were found to affect the overall calculation significantly.

This behavior was identified as a fundamental deficiency of the new model: in order to improve on it, it was chosen to move to a full-CFD description of the primary circuit including the pumps (fig. 9, right), and to use an improved fixed-point algorithm (Sec. 3.2) to couple the CATHARE pump model to the CFD. In this model, the STH primary side is completely overlapped, so that the stability/robustness concerns of the previous model are no longer present: this allows for more accurate predictions of the natural convection flowrate.

However, the move to a complete-primary CFD also incurred significant numerical costs : although the coupling iteration counts of the new model are actually lower than the intermediate model (thanks to the use of the improved fixed-point algorithms of Section 3.2), meshing the pumps, diagrid and pump-diagrid lines leads to an $\sim$20% overall mesh. Additionally, some regions (like the bottom of the cold pool) are devoid of 3D effects: however, they still must be meshed in the full-CFD model.

Hence, the trade-offs underlying the choice of a coupling strategy (choices of domain, decomposition vs overlapping) and the associated algorithms may shift when considering new application cases: the fixed-point domain-overlapping

approach was found lacking when considering a case where the overlapped and non-overlapped parts of the primary circuit became too unbalanced. To obtain a more robust coupling in natural convection, it was necessary to replace it with a complete overlap by a CFD domain, along with an improved fixed-point coupling at the primary pumps.

5. Conclusion

Although multi-scale simulations have been successfully applied in a research context to a variety of cases, from small- and medium-scale experiments to complete reactor simulations, their application to a broader range of contexts has been impeded by their high computational cost. Although some of this cost is inherent in the use of CFD, a judicious choice of coupling algorithms can be instrumental in avoiding additional overhead due to stability constraints on time steps or excessive iteration counts at each time step. On that topic, the state of the art could be summarized as follows:

1. For single-phase flow problems, a domain-overlapping approach associated with a fixed-point iteration over pressure differences (Sec. 3.1), associated with improved algorithms for thermal boundaries (Sec. 3.2), can deliver satisfactory mass, energy and momentum conservation properties in as low as 2 STH + 1 CFD iteration per timestep, thus yielding a very low coupling overhead. This approach is used by the MATHYS coupling tool developed at CEA for SFRs[15]. On the other hand, the domain-decomposition approach leads to high iteration counts: mitigation via more advanced iterative techniques (Sec. 3.3) is possible, but complex. In the absence of hydraulic coupling boundaries (for instance if the entire primary circuit is to be modelled in CFD), the decomposition and overlapping approaches are equally viable.
2. For two-phase flow problems, the domain-overlapping approach is no longer valid, making the domain-decomposition approach unavoidable. In that case, implementing a quasi- or genuine Newton method (Sec. 3.3 and 3.4)

seems to be the surest way to bring its iteration count down, as its convergence properties should be equivalent to those of a single, monolithic solution algorithm.

Apart from these considerations, the domain-overlapping approach also enables one to use the same STH input deck for both standalone and coupled calculations: this can help in comparing the differences between the system and multi-scale calculations. It can also assist in providing a "graduated" safety justification, where most transients are analyzed at the system scale (if no influent complex phenomena occur) and multi-scale is only used when necessary.

Finally, it should be noted that using multi-scale calculations in support of a safety demonstration brings in an additional set of requirements :

1. the coupled effects predicted by multi-scale simulations should be validated on a dedicated validation database, relying, as for system codes, on sets of separate-effects tests (SETs), integral-effect tests (IETs) and integral tests (SITs);
2. in order to guarantee that simulations of these tests will actually contribute to validating the final reactor application, the implementation of the coupling algorithm used should be made generic enough to apply both to the validation database and to the reactor case : this excludes, for instance, "ad-hoc" methodologies where the coupling algorithm is decided on a case-by-case basis.

Thanks to advances in two-phase CFD and in numerical performance, multi-scale simulations are certain to extend to a broader range of applications in the near future, including eventually to safety assessments of reactor transients. In the long term, their capabilities are likely to be included in a new generation of codes capable of combining natively the system, subchannel and CFD thermal-hydraulic modelling scales together with extended multi-physics modelling capabilities.

Nomenclature

CFD	Computational Fluid Dynamics
LOCA	Loss of Coolant Accident
LMR	Liquid-Metal-cooled Reactor
PWR	Pressurized Water Reactor
SGTR	Steam Generator Tube Rupture
SLB	Steam Line Break
STH	Sytem Thermal Hydraulics
SFR	Sodium Fast Reactor

References

References

[1] H. Yoon, I. Park, J. Lee, Y. Cho, S. Lee, Multi-scale and multi-physics nuclear reactor simulation for the next generation LWR safety analysis, NURETH18 (paper 6026), 2019.

[2] A. Gerschenfeld, Multiscale and multiphysics simulation of sodium fast reactors: from model development to safety demonstration, NURETH18 (paper 4164), 2019.

[3] D. Pialla, D. Tenchine, S. Li, P. Gauthe, A. Vasile, R. Baviere, N. Tauveron, F. Perdu, L. Maas, F. Cocheme, K. Huber, X. Cheng, Overview of the system alone and system/cfd coupled calculations of the phenix natural circulation test within the thins project, Nuclear Engineering and Design 290 (Supplement C) (2015) 78 – 86.

[4] A. Papukchiev, M. Jeltsov, C. Geffray, K. Kööp, P. Kudinov, R. Macian-Juan, G. Lerchl, Prediction of complex thermal-hydraulic phenomena supplemented by uncertainty analysis with advanced multiscale approaches for the tall-3d t01 experiment, 2014.

[5] K. Zwijsen, D. Martelli, P. Breijder, N. Forgione, F. Roelofs, Multi-scale modelling of the circe-hero facility, Nuclear Engineering and Design 355 (2019) 110344.

[6] H. Uitslag-Doolaard, F. Alcaro, F. Roelofs, X. Wang, A. Kraus, A. Brunett, J. Thomas, C. Geffray, A. Gerschenfeld, Multiscale modelling of the PHENIX dissymmetric test benchmark, Nuclear Engineering and Design 356 (2020) 110375.

[7] D. Grishchenko, A. Papukchiev, C. Liu, C. Geffray, M. Polidori, K. Kööp, M. Jeltsov, P. Kudinov, Tall-3d open and blind benchmark on natural circulation instability, Nuclear Engineering and Design 358 (2020) 110386.

[8] M. Saez, S. Menou, B. Uzu, The pre-conceptual design of the nuclear island of astrid, ICAPP'12, 2012.

[9] P. Turinsky, D. Kothe, Modeling and simulation challenges pursued by the consortium for advanced simulation of light water reactors (CASL), Journal of Computational Physics 313 (2016) 367–376.

[10] C. Chauliac, J.-M. Aragonés, D. Bestion, D. G. Cacuci, N. Crouzet, F.-P. Weiss, M. A. Zimmermann, Nuresim – a european simulation platform for nuclear reactor safety: Multi-scale and multi-physics calculations, sensitivity and uncertainty analysis, Nuclear Engineering and Design 241 (9) (2011) 3416 – 3426.

[11] A. Toti, J. Vierendeels, F. Belloni, Improved numerical algorithm and experimental validation of a system thermal-hydraulic/cfd coupling method for multi-scale transient simulations of pool-type reactors, Annals of Nuclear Energy 103 (2017) 36 – 48.

[12] A. Toti, J. Vierendeels, F. Belloni, Coupled system thermal-hydraulic/cfd analysis of a protected loss of flow transient in the myrrha reactor, Annals of Nuclear Energy 118 (2018) 199 – 211.

[13] I. Park, J. Lee, S. Lee, H. Yoon, J. Jeong, An implicit code coupling of 1-d system code and 3-d in-house cfd code for multi-scaled simulations of nuclear reactor transients, Annals of Nuclear Energy 59 (2013) 80 – 91.

[14] H.-M. Prasser, G. Grunwald, T. Höhne, S. Kliem, U. Rohde, F.-P. Weiss, Coolant mixing in a pressurized water reactor: Deboration transients, steam-line breaks, and emergency core cooling injection, Nuclear Technology 143 (1) (2003) 37–56.

[15] A. Gerschenfeld, Y. Gorsse, S. Li, R. Lavastre, Development and validation of multi-scale thermal-hydraulics calculation schemes for sfr applications at CEA, FR17, 2017.