https://doi.org/10.5194/egusphere-2026-1661
Preprint. Discussion started: 17 June 2026

Quantifying the scaling penalty of sequential coupling: The sequential coupling cost (C_{seq}) metric for Earth System Model components

Oriol Tintó Prims[1], Sergi Palomas[1], Simone Vacondio[1], and Mario C. Acosta[1]
[1]Barcelona Supercomputing Center, Barcelona, 08034, Spain
Correspondence: Oriol Tintó Prims (oriol.tinto@bsc.es)

Abstract. Coupling strategies in Earth System Models significantly influence their computational performance and resource efficiency. While coupling costs are traditionally evaluated in the context of concurrent implementations, the cost implications of sequential coupling approaches remain poorly quantified. In this work, we define and formalize the sequential coupling cost (C_{seq}): the additional computational resources required for a sequentially coupled ESM to achieve performance parity with an optimized concurrent setup.

We develop an analytical framework that utilizes individual component scalability profiles, derived directly from model execution logs, to diagnose inefficiencies inherent in sequential coupling architectures. To demonstrate the versatility of the metric, we apply it to three structurally distinct coupling scenarios, using models widely used by the community: atmosphere–ocean (IFS–NEMO), ocean–sea ice (NEMO–SI3), and atmosphere–aerosols (IFS–M7).

Our analysis reveals that sequential coupling imposes a substantial, yet often overlooked, efficiency deficit. By forcing all components to share a fixed resource pool, this approach ignores potential component-level parallelism and creates an exclusive reliance on domain-decomposition as a scaling strategy. While scaling solely via domain decomposition is sustainable in linear scaling regimes, it accelerates the loss of efficiency as soon as components enter sub-linear scaling regimes. We demonstrate that to achieve target throughputs, sequential setups often necessitate significant over-provisioning of computational resources, leading inefficiencies that traditional metrics fail to capture. The proposed C_{seq} metric quantifies these structural overheads, offering a quantitative basis for informing the design and architectural choices of next-generation exascale Earth System Models.

1 Introduction

Coupled Earth System Models (ESMs) are essential tools for simulating complex interactions across the climate system, including the atmosphere, ocean, land surface, and cryosphere. As scientific demands shift toward higher spatial resolution and more comprehensive process representation, the computational burden of these models increases rapidly. A twofold increase in horizontal resolution typically results in an eightfold increase in computational cost (Palmer, 2020), while the addition of complex sub-components—such as aerosols, ocean biogeochemistry, or sea-ice dynamics—introduces further challenges. Acosta

https://doi.org/10.5194/egusphere-2026-1661
Preprint. Discussion started: 17 June 2026

et al. (2024) performance study involving multiple Coupled Model Intercomparison Project Phase 6 (CMIP6) experiments showed that the coupling cost adds a computational overhead (in core-hours) of 13% in average.

Traditionally, ESMs follow one of two coupling strategies (Mechoso et al., 2021): concurrent coupling (via an MPMD approach) or sequential coupling (via a single-binary architecture). In the former, component source codes run as separate executables simultaneously to enable macro-task parallelism; in the latter, they execute one after another within a shared context. While these pairings are the norm, exceptions exist—such as single-binary architectures achieving concurrency through rank-based task partitioning, or MPMD models running sequentially. Furthermore, some frameworks do not classify function-based execution within a single context as coupling, viewing it instead as an internal sequential workflow.

Most performance metrics used today, such as those proposed by the Computational Performance Model Intercomparison Project (CPMIP) (Balaji et al., 2017), provide a broad standard for model throughput. However, their specific definition of 'coupling cost' was designed only for concurrent architectures. In these frameworks, the coupling cost is typically quantified as the component synchronization overhead or idle time between independently running components. While state-of-the-art couplers—such as OASIS (Craig et al., 2017) or YAC (Hanke et al., 2016) — provide diagnostic tools to quantify these synchronization costs in concurrent systems, they offer no mechanism to evaluate the computational overheads inherent to sequential couplings.

While the majority of ESM models utilize concurrent approaches, others remain bound to a single binary. Even within concurrent configurations, there is unexploited potential for further decoupling. For instance, in systems like EC-Earth (IFS() coupled to NEMO (Madec et al., 2024)), the sea-ice component (SI3 (Vancoppenolle et al., 2023)) remains bundled within the ocean model execution context. These models could further separate the sea-ice component to run all three elements concurrently.

However, sequential coupling strategies — prevalent in configurations like the IFS–NEMO bundle from the European Centre for Medium-Range Weather Forecasts (ECMWF)— are not captured by existing coupling cost metrics. In these setups, inefficiencies do not arise from idle time due to component synchronization, but rather from a "hidden" resource cost: by ignoring potential component-level parallelism, the scaling strategy is forced to rely exclusively on domain-decomposition and its inherent overhead. This limitation is particularly critical given the trajectory of modern HPC architectures. To fully exploit the massive concurrency of many-core processors and heterogeneous systems, model designs must maximize parallelism at all available levels. Given that components share the same pool of processes, resources cannot be allocated to match each component's specific scalability profile. Achieving a target throughput thus often requires over-provisioning for the less-scalable component—pushing it deeper into its strong-scaling limit—resulting in significant core-hour waste that traditional metrics fail to account for. Furthermore, the choice of coupling strategy can dictate the underlying numerical algorithms. For instance, enabling concurrency may necessitate different schemes or lagged exchanges that influence model accuracy. Having a standardized way to quantify the computational cost is therefore essential for evaluating the inherent trade-offs between architectural performance and physical fidelity.

https://doi.org/10.5194/egusphere-2026-1661
Preprint. Discussion started: 17 June 2026

To address this need and quantify these specific inefficiencies, we introduce the sequential coupling cost (C_{seq}), defined as the surplus computational resources required by a sequentially coupled model to match the throughput of an optimized concurrent configuration.

To demonstrate the broad applicability and usefulness of C_{seq} as a diagnostic tool, we evaluate three structurally distinct coupling scenarios representative of contemporary ESM architectures. These include the ECMWF IFS–NEMO atmosphere–ocean, the NEMO–SI3 ocean–sea ice (Madec et al., 2024; Vancoppenolle et al., 2023), and the OpenIFS–M7 atmosphere–aerosol (Huijnen et al., 2022; Vignati et al., 2004) configurations.

The continued prevalence of these sequential configurations in state-of-the-art models underscores the necessity of the C_{seq} metric, as it provides developers with the clarity needed to fully understand the performance consequences of their architectural design choices.

By extending the current suite of ESM performance metrics with C_{seq}, we provide model developers and with a standardized formulation to recognize and quantify the computational impact of single-binary coupling architectures. This provides a critical diagnostic for evaluating software architecture choices and understanding their performance implications, ensuring that design decisions are optimized to maximize model performance on next-generation exascale platforms.

2 Conceptual Illustration of Sequential Coupling Penalties

To illustrate why sequential coupling is often slower, imagine a model that scales linearly: doubling the resources to $2N$ exactly halves the execution time. In this ideal case, running two instances one after the other using $2N$ resources takes the same amount of time as running a single instance once with N resources. Similarly, running those two instances concurrently side-by-side using N resources each would also take that same amount of time. However, Earth science models typically scale sub-linearly, meaning that doubling the resources to $2N$ does not actually halve the runtime. In a sequential setup, running two instances using $2N$ resources will therefore take more time than a single run on N resources. In contrast, a concurrent setup allows each instance to run on N resources and maintain its original runtime. This difference makes the concurrent approach faster for the same total resource budget. This can be seen in Figure 1

https://doi.org/10.5194/egusphere-2026-1661
Preprint. Discussion started: 17 June 2026

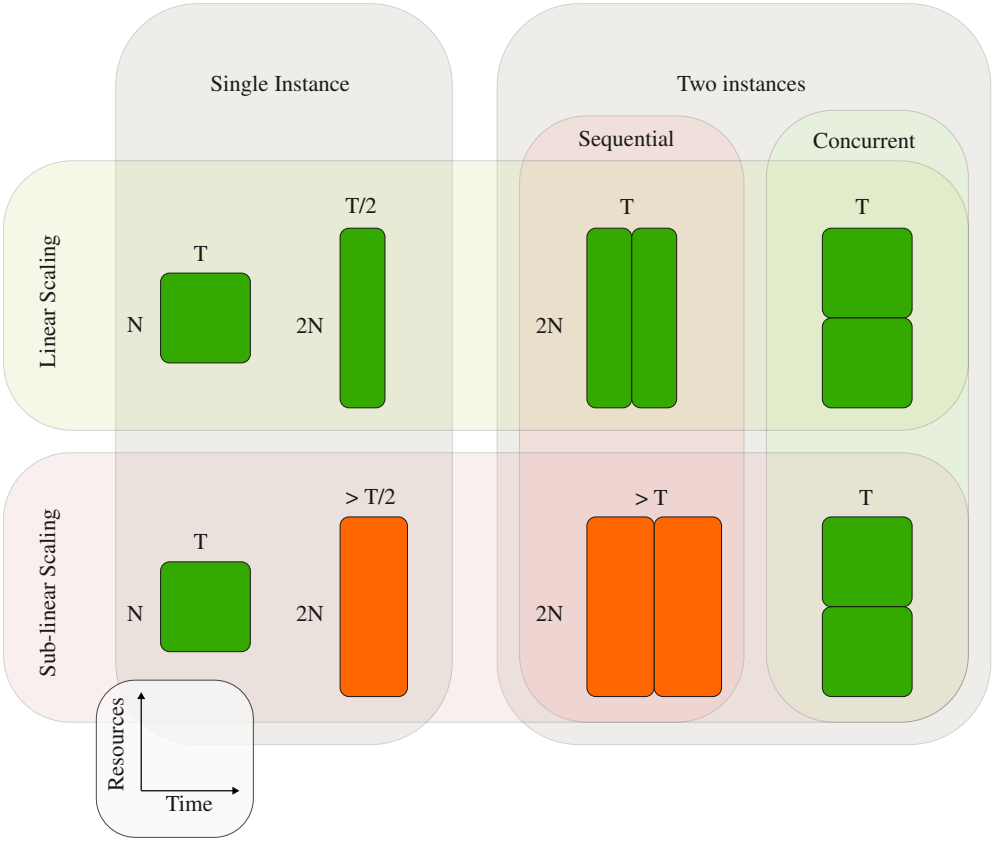

Figure 1. Comparison of total wall-clock time for different coupling strategies. While concurrent coupling maintains the baseline runtime when doubling the resources, when the models exhibit sub-linear scaling, the sequential setup results in a greater total execution time.

https://doi.org/10.5194/egusphere-2026-1661
Preprint. Discussion started: 17 June 2026

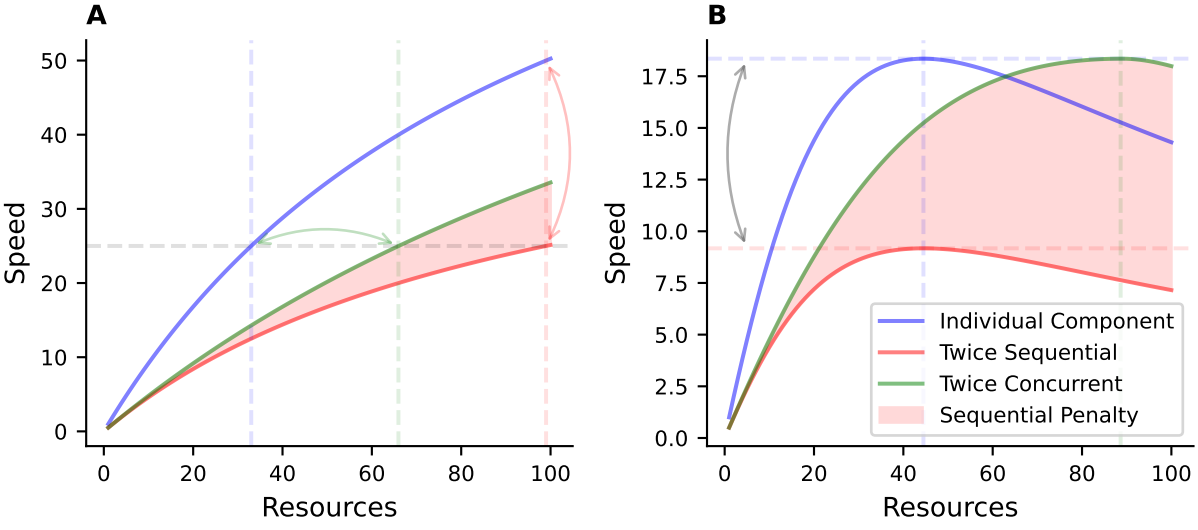

Figure 2. Scaling performance and resource efficiency comparison between sequential and concurrent coupling configurations. The shaded red areas highlight the speed and resource penalties incurred by the sequential coupling. (A) Case A illustrates the resource overhead triggered by sub-linear scaling; while the concurrent approach (green) requires exactly twice the resources ($2N$) to maintain the speed of a single component (N; see green arrow), the sequential setup (red) must use significantly more to reach that same speed, with the resource gap widening as the component's scaling efficiency degrades. In this regard, it can be seen that, at fixed resources, the speed of the sequential approach is half of that of the individual component (red arrow). (B) Case B demonstrates a hard performance ceiling where the sequential configuration reaches a peak speed exactly half that of the concurrent case, represented by the gray arrow.

As illustrated in Figure 2 A, a sequential setup results in a total completion time that is twice the runtime of the individual component, whereas the concurrent setup completes in the time required for a single instance using half of the total resources. This distinction highlights how a sequential configuration fundamentally limits achievable speed. As shown in Figure 2 B, because the sequential approach doubles the total runtime, the maximum speed is effectively halved. Conversely, a concurrent setup maintains the speed of the single component by simply doubling the resource allocation, ensuring that the system capacity scales effectively with the hardware.

3 Defining Sequential Coupling Cost

In (Balaji et al., 2017) the coupling cost was defined as:

$$C = \frac{T_M P_M - \sum_c T_c P_c}{T_M P_M} \tag{1}$$

https://doi.org/10.5194/egusphere-2026-1661
Preprint. Discussion started: 17 June 2026

where T_M and P_M are the runtime and parallelization for the whole model, and T_c and P_c the same for individual components.

However, since in a sequential case $T_M P_M = \sum_c T_c P_c$ this results in the C being always 0. As was shown in the previous section, in sub-linear scaling regimes concurrent approaches outperform sequential ones, and the coupling cost metric fails to capture this reality.

To define a metric that can successfully quantify the sequential coupling cost (C_{seq}, let us consider two components A and B. Being $T_A(N_A)$ the time it takes for A to complete a step when using N_A processes and $T_B(N_B)$ the time it takes for B to complete a step when using N_B processes. In a sequential setup, the two components are executed one after the other using the same resources, so

$$T_{\text{seq}}(N) = T_A(N) + T_B(N) \tag{2}$$

And the throughput of the sequential scheme (S_{seq}) can be computed as the inverse of time like:

$$S_{\text{seq}}(N) = \left(\frac{1}{S_A(N)} + \frac{1}{S_B(N)} \right)^{-1} \tag{3}$$

In the concurrent setup, components A and B are executed simultaneously on disjoint resource partitions N_A and N_B, such that $N_A + N_B = N$. Therefore, the time it takes for running components A and B in a concurrent coupling scheme (T_{conc}) is

$$T_{\text{conc}}(N) = \max\left(T_A(N_A), T_B(N_B)\right)\Big|_{N_A+N_B=N} \tag{4}$$

leading to a concurrent throughput (S_{conc})

$$S_{\text{conc}}(N) = \min(S_A(N_A), S_B(N_B))\Big|_{N_A+N_B=N} \tag{5}$$

Let us assume that the function $S(N)$ can be inverted to obtain $N(S)$.

For the sequential case, this would be $N_{\text{seq}}(S)$, for the concurrent case $N_{\text{conc}}(S) = N_A(S) + N_B(S)$.

To have a metric that can capture the resource overhead needed due to a sequential coupling, we propose, for a target speed S^*:

$$C_{\text{seq}}(S^*) = 100 \cdot \frac{N_{\text{seq}}(S^*) - N_{\text{conc}}(S^*)}{N_{\text{seq}}(S^*)} = 100 \cdot \frac{N_{\text{seq}}(S^*) - N_A(S^*) - N_B(S^*)}{N_{\text{seq}}(S^*)} \tag{6}$$

So the resources that are needed to reach a target throughput in a sequential setup, minus the resources that are needed to reach the same target in a concurrent setup, normalized.

This represents the percentage of the resources that are needed solely because of the inefficiencies of the sequential coupling.

https://doi.org/10.5194/egusphere-2026-1661
Preprint. Discussion started: 17 June 2026

Measuring traditional coupling costs typically requires only a single execution. In contrast, computing C_{seq} requires us to model the performance of individual components to estimate their run-times under different resource configurations. Provided we can extract isolated execution data for each component to formulate an accurate analytical function, C_{seq} can be successfully derived.

4 Practical Demonstrations

To validate the proposed sequential coupling cost (C_{seq}), we apply the metric to three structurally distinct Earth System Model (ESM) configurations. These scenarios represent different types of physical interactions commonly evaluated by the community: atmosphere–ocean (IFS–NEMO), ocean–sea ice (NEMO–SI3), and atmosphere–aerosols (OpenIFS–M7). By computing C_{seq} for these diverse cases, we demonstrate the metric's broad applicability in identifying structural bottlenecks.

The selected configurations rely on widely used, state-of-the-art components. The Integrated Forecasting System (IFS), developed by the European Centre for Medium-Range Weather Forecasts (ECMWF), serves as the atmospheric component. Ocean dynamics are handled by NEMO (Nucleus for European Modelling of the Ocean) (Madec et al., 2024), while its integrated sea ice engine, SI3 (Sea Ice modelling Integrated Initiative) (Vancoppenolle et al., 2023), drives the ocean–sea ice scenario. Finally, aerosol interactions are represented using the M7 microphysics model (Huijnen et al., 2022) coupled to OpenIFS (Sparrow et al., 2021). All simulations presented in this section were executed on the general-purpose partition of the MareNostrum 5 supercomputer, hosted by the Barcelona Supercomputing Center (BSC).

A key practical advantage of this metric is its reliance on existing diagnostic outputs rather than external profiling tools. To compute the scalability profiles and the resulting C_{seq}, we extract component run-times directly from the standard execution logs inherently generated by the models. Because these models already maintain basic temporal records, applying our metric requires no complex code instrumentation or specialized performance analysis software. However, the format and granularity of these logs naturally vary across different ESMs. While some models natively output clearly delineated component times, others necessitate a degree of post-processing to accurately isolate and attribute the integration times to their respective components.

Crucially, our timing measurements strictly isolate the stepping time —the duration of the main integration loop— and exclude all model initialization and finalization phases. Coupling imbalances and sequential overheads primarily dictate the throughput of the recurring execution phase; therefore, isolating the stepping time provides the most accurate measure of coupling efficiency. Using this empirical stepping data extracted from the logs, we construct the analytical scalability functions for each component. To standardize and automate this analytical workflow, all scalability profiles and subsequent metric evaluations are computed using the *Escalador* library (Tintó Prims and Palomas, 2026). These functions then allow us to compute the theoretical run-times under different resource allocations, enabling the derivation of the optimal configurations and the final C_{seq} reported for each scenario.

https://doi.org/10.5194/egusphere-2026-1661
Preprint. Discussion started: 17 June 2026

4.1 Use Case: IFS–NEMO (Atmosphere–Ocean)

The IFS–NEMO configuration used by ECMWF and within the DestinE (Wedi et al., 2022; Doblas-Reyes et al., 2025) project employs a sequential strategy where the NEMO ocean model is integrated via a function call inside the IFS atmospheric executable, first described in Mogensen et al. (2012). The specific experimental setup and component resolutions are summarized in Table 1.

Table 1. Experimental configuration for the IFS–NEMO sequential coupling use case.

Component	Type	Version	Grid	Resolution	Horiz. Points	Levels	Total 3D Points
IFS	Atmosphere	48r1	Tco2559	∼4.4 km	65,996,800	137	$\sim 9.04 \times 10^9$
NEMO	Ocean	4.0	eORCA12	$1/12^\circ$	13,221,000	75	$\sim 0.99 \times 10^9$

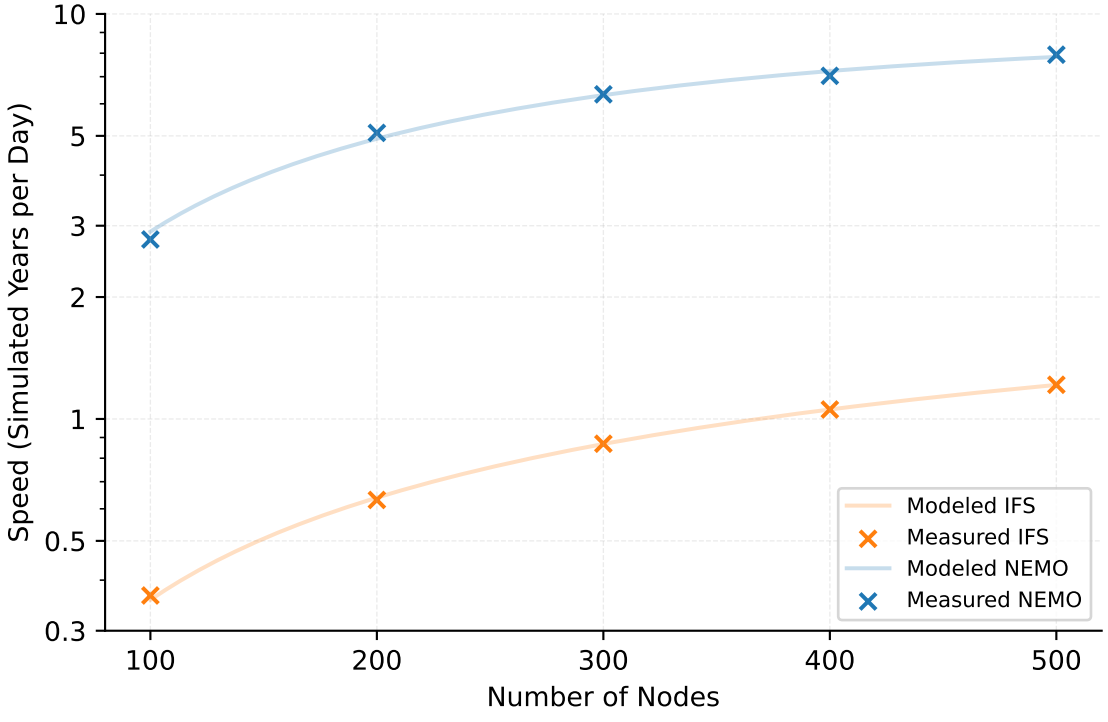

Figure 3. Scalability of IFS and NEMO components. Crosses indicate measured performance from execution logs; solid lines represent the analytical model fits. Speed is expressed in simulated years per day (SYPD) as a function of node count.

For this evaluation, timing data were extracted from standard execution logs across varying node counts (N) to characterize individual component scalability. To model performance, a Universal Scalability Law (USL)(Gunther, 2008) was fitted to the throughput data. As illustrated in Figure 3, a high-fidelity representation of observed performance is provided by these ana-

https://doi.org/10.5194/egusphere-2026-1661
Preprint. Discussion started: 17 June 2026

lytical fits, and the saturation point—where NEMO's performance degrades due to communication contention—is accurately captured.

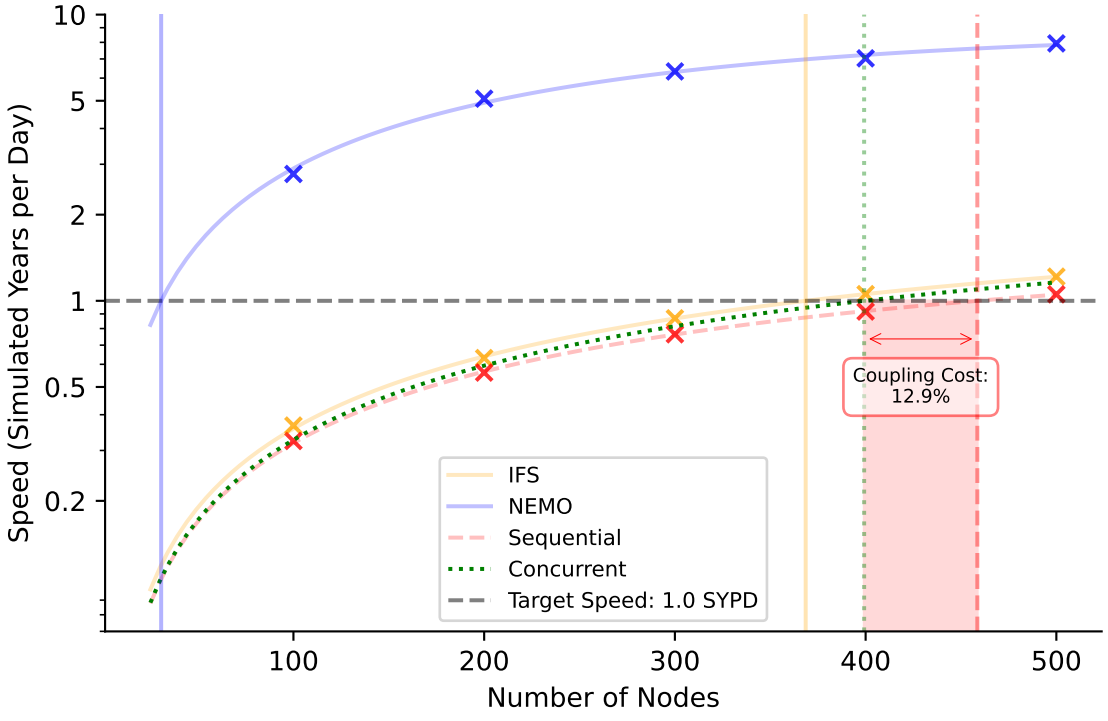

Figure 4. Scalability of IFS (orange) and NEMO (blue) components as well as the sequentially coupled model (red) and the concurrently coupled model (green). Speed is expressed in simulated years per day (SYPD) as a function of node count. The shaded area represents the resource overhead of the sequential coupling against its concurrent counterpart.

As can be seen in Figure 4, the computational requirements were evaluated using these analytical functions for a target throughput of $S^* = 1$ Simulated Year Per Day (SYPD), a key performance milestone for DestinE (Segura et al., 2025). In the sequential configuration, $N_{\text{seq}} = 458.6$ nodes are required by the integrated bundle to achieve this target. However, it is revealed by independent scalability models that 368.6 nodes are required for the atmosphere and only 30.6 nodes are required for the ocean to meet the same throughput. Thus, the optimal concurrent requirement is defined as:

$$N_{\text{conc}}(1\ \text{SYPD}) = N_{\text{IFS}}(1\text{SYPD}) + N_{\text{NEMO}}(1\text{SYPD}) = \underbrace{368.6\ \text{nodes}}_{\text{IFS}} + \underbrace{30.6\ \text{nodes}}_{\text{NEMO}} = 399.2\ \text{nodes} \tag{7}$$

Applying the sequential coupling cost metric (C_{seq}), we find:

https://doi.org/10.5194/egusphere-2026-1661
Preprint. Discussion started: 17 June 2026

$$C_{\text{seq}}(1\ \text{SYPD}) = 100 \cdot \frac{458.6 - 399.2}{458.6} = 12.9\% \tag{8}$$

This C_{seq} value of 12.9% indicates that nearly one-eighth of the total computational resources are lost due to the sequential architecture. In this setup, NEMO is forced to run on 458.6 nodes—nearly fifteen times the 30.6 nodes actually needed to reach 1 SYPD. This over-provisioning pushes the ocean component deep into an inefficient regime where communication overhead dominates. As a result, the total runtime is no longer governed by calculation speed, but by the structural constraints of the coupling. In contrast, a concurrent approach would allow NEMO to operate with much higher efficiency while the atmosphere scales independently, achieving the same throughput with significantly lower core-hour consumption. Crucially, this structural shift restores the value of model-level optimization; in a concurrent regime, algorithmic refinements translate directly into reduced resource usage rather than being masked by communication overhead.

The impact of this structural inefficiency is clearly reflected in the final resource budget. Analysis of computational costs shows that NEMO is nearly twice as expensive to run in a sequential configuration: costs rise from $82.4\ \text{k}\cdot\text{core}\cdot\text{hour/year}$ (30.6 nodes) in the concurrent setup to $161.9\ \text{k}\cdot\text{core}\cdot\text{hour/year}$ (458.6 nodes) in the sequential setup, representing a 96% surge in resource allocation.

Similarly, the IFS component scales from $990.7\ \text{k}\cdot\text{core}\cdot\text{hour/year}$ (399.2 nodes) to $1070.8\ \text{k}\cdot\text{core}\cdot\text{hour/year}$ (458.6 nodes). Although this represents a modest 8% relative increase, the larger computational footprint of the atmospheric component results in an absolute increment of $80.1\ \text{k}\cdot\text{core}\cdot\text{hour/year}$. Notably, this growth in resource demand is nearly identical to the total cost increase observed in NEMO, illustrating that even small percentage shifts in the atmospheric component carry significant computational weight.

It should be noted that the ideal node count for the ocean component (30.6) falls below our empirical range of 100–500 nodes. While this extrapolation introduces some uncertainty, the use of the USL equation likely keeps the deviation from actual performance within a small margin.

4.2 Use Case: NEMO–SI3 (Ocean–Sea-ice)

According to literature (Tintó Prims et al., 2019; Irrmann et al., 2022), the main obstacle to NEMO's scalability is the communication overhead inherent in domain decomposition. While these studies also present work aimed at alleviating this issue, the fundamental limitation persists.

Previous studies have explored decoupling sea-ice processes from the core ocean computations within NEMO to enable concurrent execution and improve load balancing (Maisonnave and Masson, 2015; Maisonnave and Bourdallé-Badie, 2022).

Given these considerations, an analytical assessment of the ocean–sea-ice coupling cost can identify the most effective coupling strategy. The experimental setup and component specifications are summarized in Table 2.

Component-level modeling (Figure 5) reveals a significant performance disparity: while the sea-ice component is inherently faster than the ocean component, it exhibits considerably poorer scaling, which can be understood from the reduced computational cost and the relatively higher demand for communication. While concurrent execution is generally more efficient if

https://doi.org/10.5194/egusphere-2026-1661
Preprint. Discussion started: 17 June 2026

Table 2. Experimental configuration for the NEMO–SI3 sequential coupling use case.

Component	Type	Version	Grid	Resolution (Eq.)	Horiz. Points	Levels	Total 3D Points
NEMO	Ocean	5.0	ORCA025	∼28 km	1,472,282	75 Levels	$\sim 1.10 \times 10^8$
SI3	Sea-ice	5.0	ORCA025	∼28 km	1,472,282	5 Categories	$\sim 7.36 \times 10^6$

individual components exhibit sub-linear scaling, the performance disparity between the ocean and sea-ice further justifies this approach. By utilizing concurrent coupling, each component can operate within a more efficient resource regime. In this specific scenario, a concurrent strategy allows the fast, poorly scaling sea-ice component to be restricted to a small fraction of resources, while the majority are devoted to the slower, more scalable ocean component.

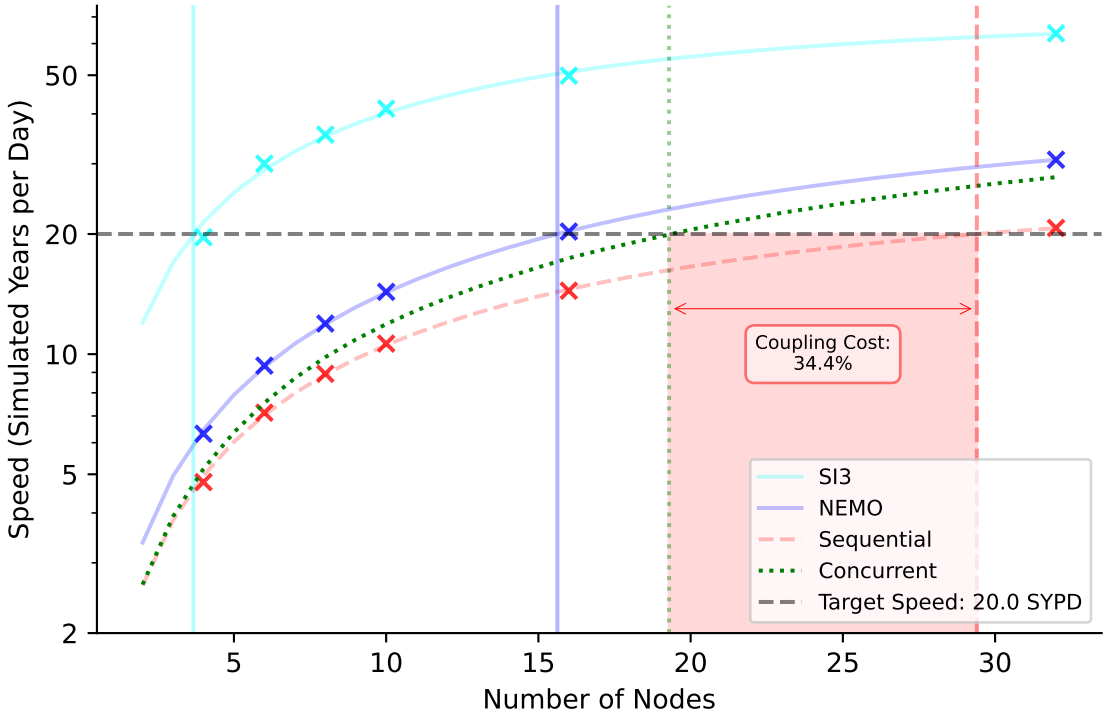

Figure 5. Analytical performance models for ORCA025. Crosses indicate empirical data. The dotted green line represents the optimal concurrent model derived from the analytical models. The sea-ice component (cyan) is faster but saturates earlier than the ocean component (green). The ideal concurrent coupling outperforms the sequential coupling leading to a 34.4% sequential coupling cost for a target speed of 20 SYPD.

Starting from the run as 4 nodes as the reference, the ocean component has an 80% efficiency at 20 Simulated Years per Day. While efficiency requirements can also be arbitrary, we consider that this is a good speed target to compute our sequential coupling cost C_{seq}.

https://doi.org/10.5194/egusphere-2026-1661
Preprint. Discussion started: 17 June 2026

In the current sequential configuration, the sea-ice component is forced to utilize the same number of nodes (N) as the ocean, leading to a rapid escalation in resource overhead as the system scales. Based on our analytical models, our components can reach 20 SYPD with at $N_{\text{NEMO}} = 15.6$ nodes and $N_{\text{SI3}} = 3.7$, leading to a $N_{\text{conc}} = N_{\text{NEMO}} + N_{\text{SI3}} = 15.6 + 3.7 = 19.3$. On the other hand the sequential setup needs $N_{\text{seq}} = 29.4$. Therefore, using equation 6, the sequential coupling cost becomes $C_{\text{seq}} = 34.38\%$.

Similarly to the other use cases, C_{seq} increases with the total number of nodes N.

The structural inefficiency observed in the previous use case is even more pronounced here, as reflected in the final resource budget. Similar to the IFS-NEMO configuration, the secondary component—in this case, SI3—suffers the most from sequential constraints. Its costs jump from 492.9 core·h/year (3.7 nodes) to 1.3k core·h/year (29.4 nodes), representing a 157.2% surge that more than doubles its original allocation.

The NEMO component follows a similar trend to the IFS atmospheric component, though with a more significant relative impact. It scales from 2.1k core·h/year (15.6 nodes) to 2.7k core·h/year (29.4 nodes). While this 27.8% increase is numerically larger than the 8% shift seen in IFS, the underlying logic remains: the absolute increment in the dominant component (0.6k core·h/year) is nearly identical to the total cost increase of the smaller SI3 component.

Ultimately, the aggregate data confirms a substantial overhead for the coupling method; transitioning from concurrent to sequential execution results in a total system cost increase from 2.6k to 4.0k core·h/year, a 52.4% resource penalty.

4.3 Use Case: OpenIFS–M7 (Atmosphere–Chemistry)

The EC-Earth coupling scheme for the aerosol chemistry component has evolved from a concurrent strategy in EC-Earth3 (Van Noije et al., 2021) to a sequential coupling in the latest release, EC-Earth4. In the current OIFS–M7 configuration (atmosphere and aerosols, respectively), both components are integrated into a single binary and executed sequentially, one after the other. In EC-Earth3, aerosol chemistry was computed every 6 hours using the TM5 scheme (van Noije et al., 2014), whereas in EC-Earth4 it is computed at every time step using the M7 module. In addition, atmosphere and aerosols now share the same grid, unlike in the previous version: as a consequence, the horizontal resolution for the aerosols component has increased substantially, from approximately 300 km to 80 km at the equator. These differences in coupling strategy, grid configuration, and resolution, which are summarized in Table 3, afford unprecedented accuracy in the modeling of aerosols within EC-Earth.

Previous studies have quantified the significant computational overhead introduced by chemistry computations (Christou et al., 2016), and the coupling cost within EC-Earth3—including the EC-Earth3-AerChem configuration—has been analyzed in detail (Acosta et al., 2023), reporting a coupling cost of 8.25%. Although direct comparisons between EC-Earth3 and EC-Earth4 are not straightforward due to differences in model configuration and HPC platform, this reference provides useful context for interpreting the results presented here.

All simulations were performed using an atmospheric grid of TL255 with 91 vertical levels and a 45-minute time step. Timing data were averaged over three runs of five simulated days each. In this use case, rather than prescribing a fixed target throughput, we analyze the maximum achievable throughput under a constraint of 32 compute nodes.

https://doi.org/10.5194/egusphere-2026-1661
Preprint. Discussion started: 17 June 2026

Two chemistry configurations for EC-Earth4 were considered, differing primarily in the number of chemical species simulated: simplified chemistry (SIMC, Remy and Bock (2018)) and full chemistry (FULC, Tegen et al. (2019)). Figures 6 and 7 show that OIFS throughput is largely insensitive to the chemistry configuration. In contrast, for the M7 module, SIMC is approximately twice as fast as FULC under the same parallelisation level.

Table 3. Experimental configuration for the OIFS–M7 sequential coupling use case as implemented in EC-Earth4. A comparison is shown with the TM5 aerosol scheme, which was concurrently coupled to the IFS in EC-Earth3 (Van Noije et al. (2021)).

ESM	Component	Version	Type	Grid	Resolution (Eq.)	Horiz. Points	Levels	Total 3D Points
EC-Earth4	OIFS	48r1	Atmosphere	TL255	~ 80 km	88,838	91 Levels	$\sim 8.08 \times 10^6$
	M7		Aerosols	TL255	~ 80 km	88,838	91 Levels	$\sim 8.08 \times 10^6$
EC-Earth3	IFS	36r4	Atmosphere	TL255	~ 80 km	88,838	91 Levels	$\sim 8.08 \times 10^6$
	TM5		Aerosols	$3° \times 2°$ †	~ 300 km	21,600	34 Levels	$\sim 7.34 \times 10^5$

†(latitude × longitude)

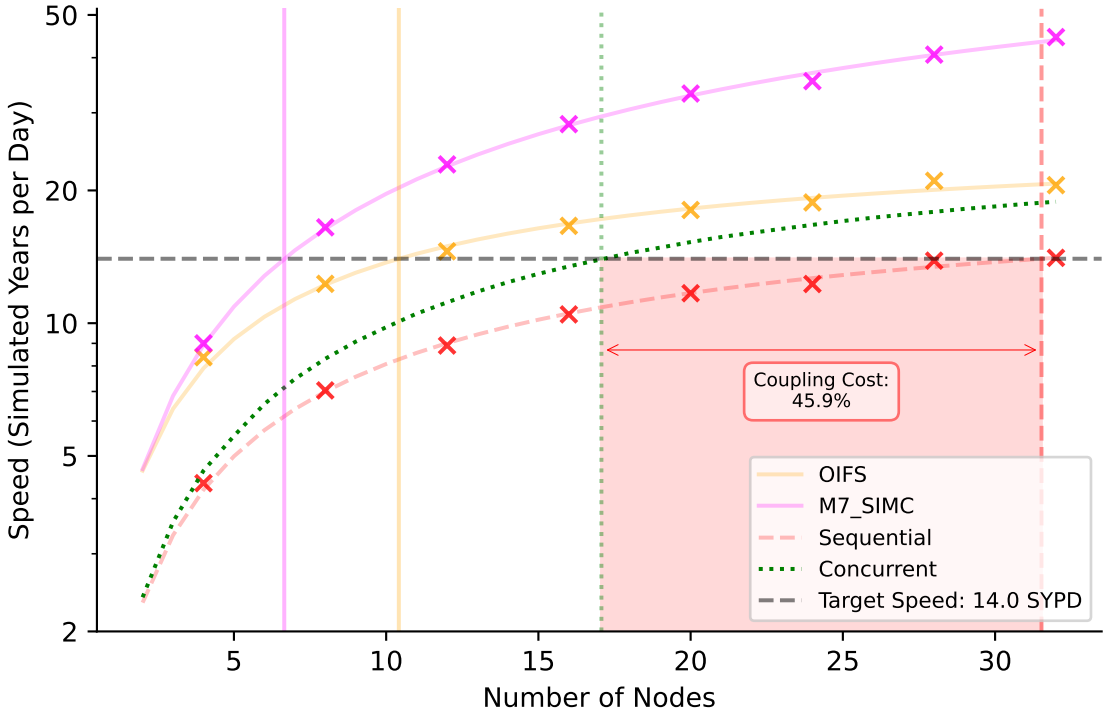

Figure 6. Resource-throughput analysis for OIFS-M7 simplified chemistry (SIMC). The horizontal line marks the 14 SYPD target, which is the maximum throughput that the sequential setup can reach with 32 nodes. The shaded region denotes the C_{seq}, illustrating the "excess" resources required by the sequential configuration compared to an optimized concurrent setup.

https://doi.org/10.5194/egusphere-2026-1661
Preprint. Discussion started: 17 June 2026

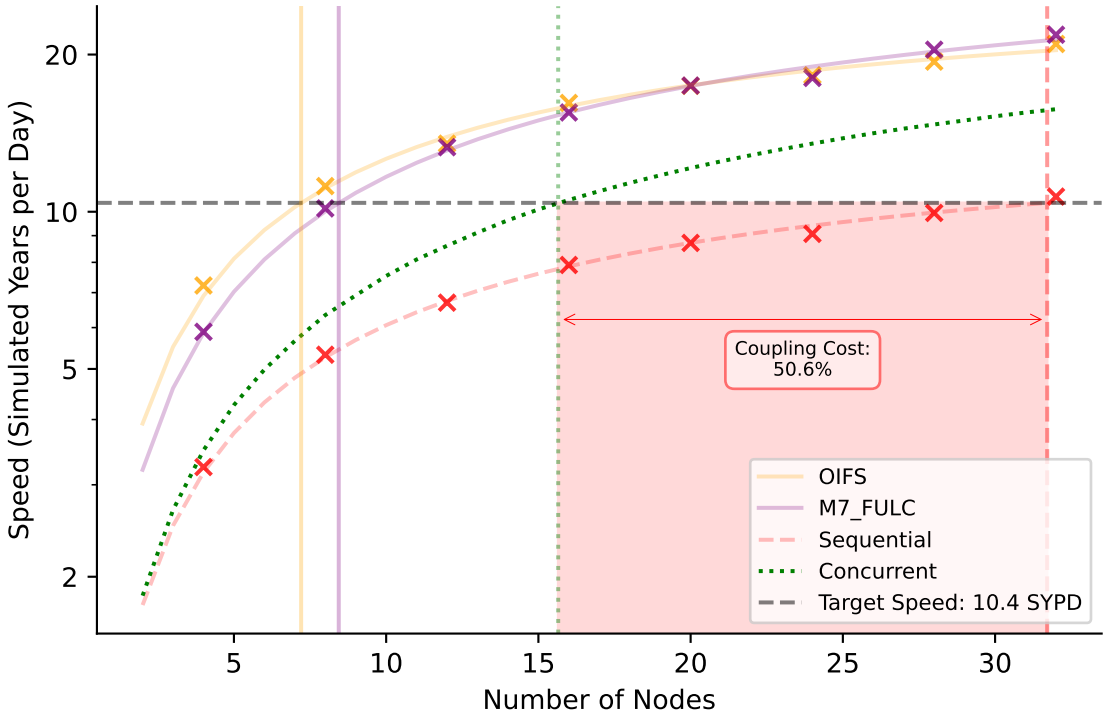

Figure 7. Resource-throughput analysis for OIFS-M7 full chemistry (FULC). The horizontal line marks the 10.4 SYPD target which is the maximum throughput that the sequential system can reach with 32 nodes. The shaded region denotes the C_{seq}, illustrating the "excess" resources required by the sequential configuration compared to an optimized concurrent setup.

Results indicate that C_{seq} becomes increasingly large at higher node counts, where the parallel efficiency of both the atmospheric and aerosol components starts to degrade. Under the 32-node constraint, C_{seq} reaches 46.2% for the SIMC configuration and 50.9% for FULC.

For the SIMC configuration, the chemistry model (M7) is significantly faster than the atmospheric model (OIFS). Under a concurrent coupling framework, we estimate that M7 would reach the target throughput of 14.1 SYPD with 7 nodes, while OIFS would require at most 11 nodes, totaling 18 nodes. Assuming a conservative concurrent coupling overhead of 10%—already above what was achieved when balancing the EC-Earth3-AerChem configuration—the same target throughput could be achieved with approximately 20 nodes.

For the FULC configuration, the throughput and scalability of OIFS and M7 are more similar, and parallel efficiency decreases rapidly as node counts increase. The maximum throughput achievable with 32 nodes is 10.4 SYPD (the target speed shown in Fig. 7). Under a concurrent coupling strategy, OIFS would reach this target with 7 nodes, while M7 would require no more than 9 nodes, totaling 16 nodes. Again, if we assume a coupling cost of 10% for the concurrent coupling scheme, the same throughput could be achieved with 17 nodes, and not with the 32 that are currently required.

https://doi.org/10.5194/egusphere-2026-1661
Preprint. Discussion started: 17 June 2026

While OIFS scalability remains largely unchanged between SIMC and FULC, M7 scales significantly worse in the FULC configuration. As a result, resource over-provisioning becomes even more pronounced, with C_{seq} reaching 50.9% under the 32-node constraint.

In EC-Earth3, these components were coupled concurrently via an external coupler, and after configuration balancing the reported coupling cost was 8.25%. Although it is not possible to perform a direct comparison between EC-Earth3 and EC-Earth4 due to differences in model formulation, resolution, and HPC systems (MareNostrum4 vs. MareNostrum5), the present results strongly suggest that the sequential coupling strategy leads to substantially higher coupling costs. A concurrent approach would very likely provide improved computational efficiency.

These findings underscore the importance of quantitative metrics such as C_{seq} for guiding coupling design decisions and supporting informed development choices within the EC-Earth consortium.

5 Discussion and Conclusions

5.1 Methodological Considerations

The analytical modeling approach using USL proved robust across diverse component types and scaling regimes. The method's reliance on empirical calibration data makes it applicable to any ESM configuration where component-level timing information can be extracted. However, several limitations warrant discussion:

Idealized Concurrent Assumptions: Our calculations assume that concurrent components could be allocated resources independently without additional coupler overhead. In practice, field exchange via external couplers (e.g., OASIS, ESMF) introduces communication costs that would partially offset the gains quantified by C_{seq}. Nevertheless, multiple studies have shown that well-optimized concurrent coupling frameworks typically incur only 1-5% overhead (Balaji et al., 2017), substantially lower than the 13-50% inefficiencies identified here.

Communication Infrastructure Dependencies: Component scalability is strongly influenced by network topology and MPI implementation characteristics. The analytical fits capture these effects implicitly through the calibration data, but transferability across platforms may require re-calibration. The good agreement between modeled and observed performance in our MareNostrum 5 experiments suggests the approach is robust within a given HPC environment.

Load Imbalance and Temporal Variability: Our analysis assumes steady-state performance characteristics. In reality, computational costs vary with simulation state (e.g., seasonal cycles, grid partitioning effects). For components with highly variable workloads, time-averaged scalability curves may underestimate peak resource requirements.

5.2 Summary of Results

Our analysis across three structurally distinct coupling scenarios demonstrates that sequential coupling cost (C_{seq}) is a pervasive and quantifiable phenomenon in modern Earth System Models. The metric reveals substantial coupling inefficiencies that have been largely overlooked.

https://doi.org/10.5194/egusphere-2026-1661
Preprint. Discussion started: 17 June 2026

The three use cases examined span a representative range of component coupling patterns:

- IFS–NEMO exhibited a C_{seq} of 13% at 1 SYPD, indicating that nearly one-eight of computational resources are wasted due to the sequential coupling. Interestingly, it can be seen than the ocean component efficiency drops by half in the sequential setup, and although the efficiency of the much heavier atmosphere component only drops by a 8%, its absolute increase in cost is of a similar magnitude than the ocean one.
- NEMO–SI3 showed a similar case in the sense that the computational requirements of both components were significantly different. Targeting a speed of 20 SYPD, for which the ocean component still maintained a decent 80% efficiency, the resulting C_{seq} was a 34.4%, showing that more than a third of the resources are effectively wasted due to the sequential coupling penalty.
- OIFS–M7 showed a case that was significantly different, because in this case the components had a much more similar cost and scaling behavior than the previous two cases. In this use-case, that covered two aerosol configurations, the coupling cost for speeds that can be reached in a 32 node allocation resulted between 40% and 50%, showing that the scalability of the model is heavily constrained by the sequential coupling.

5.3 Broader Implications for ESM Development

The sequential coupling cost metric fills a critical gap in the ESM performance evaluation toolkit. While CPMIP metrics (Balaji et al., 2017) effectively characterize synchronization overhead and load imbalance in concurrent architectures, they cannot capture the inherent efficiency losses of sequential execution within single-binary models. C_{seq} provides a standardized, quantitative basis for informed architectural decision-making.

The pursuit of numerical accuracy and software simplicity in model development typically entails a performance trade-off, resulting in a measurable impact on throughput. Until now, the lack of a formal method to quantify the computational penalty of selecting a sequential versus a concurrent coupling meant that developers lacked a critical tool to evaluate such trade-offs objectively. In this regard, the transition from EC-Earth3 to EC-Earth4 chemistry coupling serves as a clear example: while the update introduced a suite of improvements targeting higher accuracy, it also involved a structural shift to a sequential approach. As demonstrated in this study, this architectural change imposed a significant constraint on the model's scalability.

Beyond architectural design, C_{seq} serves as a decision-support tool for resource allocation. Quantifying the sequential coupling cost allows developers to assess whether the engineering effort required to implement a concurrent setup is justified by the expected performance gains. For instance, if a model is intended for low node-count executions where the difference between coupling strategies is marginal, a sequential approach may be acceptable. Conversely, if a user aims to maximize throughput and C_{seq} reveals a substantial scalability bottleneck, the metric provides the objective justification needed to invest in a concurrent configuration.

As ESMs target exascale platforms, any opportunity to exploit parallelism must be taken seriously. The sequential coupling cost metric provides a robust diagnostic to ensure that fundamental design decisions are informed by a precise, quantitative estimation of their performance implications. This transparency is essential for long-term HPC planning, ensuring that model

https://doi.org/10.5194/egusphere-2026-1661
Preprint. Discussion started: 17 June 2026

architectures do not inadvertently introduce bottlenecks that prevent them from fully leveraging next-generation computational resources.

5.4 Conclusions

Sequential coupling remains a common architectural pattern in Earth System Modeling, yet its computational costs have lacked rigorous quantification until now. We have introduced the sequential coupling cost metric (C_{seq}) as a standardized diagnostic that extends existing ESM performance evaluation frameworks to capture the inefficiencies inherent to single-binary executable designs.

Our analysis of three representative use cases—spanning atmosphere–ocean, ocean–sea ice, and atmosphere–chemistry coupling—reveals that sequential architectures routinely waste 10–50% of allocated computational resources compared to optimized concurrent alternatives. These inefficiencies arise not from component imbalance or synchronization overhead, but from a structural inability to exploit component parallelism, forcing the model to rely solely on domain decomposition and its associated overheads.

As the ESM community pursues exascale computing and increasingly complex process representation, architectural decisions that prioritize software convenience over computational efficiency—as seen in the trade-offs discussed regarding recent ESM transitions—will exact growing costs in both scientific productivity and environmental impact. However, C_{seq} also serves as a pragmatic guide for resource allocation: it allows developers to identify when the performance gain of a concurrent setup justifies the engineering investment, or when a sequential approach remains an acceptable compromise for low-resource configurations.

By integrating C_{seq} into standard performance evaluation protocols alongside existing CPMIP metrics, the community can ensure that next-generation ESM architectures are designed not only for scientific fidelity but also for computational sustainability and maximum throughput on future HPC infrastructures.

https://doi.org/10.5194/egusphere-2026-1661
Preprint. Discussion started: 17 June 2026

References

Acosta, M. C., Palomas, S., and Tourigny, E.: Balancing EC-Earth3 Improving the Performance of EC-Earth CMIP6 Configurations by Minimizing the Coupling Cost, Earth and Space Science, 10, e2023EA002 912, https://doi.org/https://doi.org/10.1029/2023EA002912, e2023EA002912 2023EA002912, 2023.

Acosta, M. C., Palomas, S., Paronuzzi Ticco, S. V., Utrera, G., Biercamp, J., Bretonniere, P.-A., Budich, R., Castrillo, M., Caubel, A., Doblas-Reyes, F., Epicoco, I., Fladrich, U., Joussaume, S., Kumar Gupta, A., Lawrence, B., Le Sager, P., Lister, G., Moine, M.-P., Rioual, J.-C., Valcke, S., Zadeh, N., and Balaji, V.: The computational and energy cost of simulation and storage for climate science: lessons from CMIP6, Geoscientific Model Development, 17, 3081–3098, https://doi.org/10.5194/gmd-17-3081-2024, 2024.

Balaji, V., Maisonnave, E., Zadeh, N., Lawrence, B. N., Biercamp, J., Fladrich, U., Aloisio, G., Benson, R., Caubel, A., Durachta, J., Foujols, M.-A., Lister, G., Mocavero, S., Underwood, S., and Wright, G.: CPMIP: measurements of real computational performance of Earth system models in CMIP6, Geoscientific Model Development, 10, 19–34, https://doi.org/10.5194/gmd-10-19-2017, 2017.

Christou, M., Christoudias, T., Morillo, J., Alvarez, D., and Merx, H.: Earth system modelling on system-level heterogeneous architectures: EMAC (version 2.42) on the Dynamical Exascale Entry Platform (DEEP), Geoscientific Model Development, 9, 3483–3491, https://doi.org/10.5194/gmd-9-3483-2016, 2016.

Craig, A., Valcke, S., and Coquart, L.: Development and performance of a new version of the OASIS coupler, OASIS3-MCT_3.0, Geoscientific Model Development, 10, 3297–3308, https://doi.org/10.5194/gmd-10-3297-2017, 2017.

Doblas-Reyes, F. J., Kontkanen, J., Sandu, I., Acosta, M., Al Turjmam, M. H., Alsina-Ferrer, I., Andrés-Martínez, M., Arriola, L., Axness, M., Batlle Martín, M., Bauer, P., Becker, T., Beltrán, D., Beyer, S., Bockelmann, H., Bretonnière, P.-A., Cabaniols, S., Caprioli, S., Castrillo, M., Chandrasekar, A., Cheedela, S., Correal, V., Danovaro, E., Davini, P., Enkovaara, J., Frauen, C., Früh, B., Gaya Àvila, A., Ghinassi, P., Ghosh, R., Ghosh, S., González, I., Grayson, K., Griffith, M., Hadade, I., Haine, C., Hartick, C., Haus, U.-U., Hearne, S., Järvinen, H., Jiménez, B., John, A., Juchem, M., Jung, T., Kegel, J., Kelbling, M., Keller, K., Kinoshita, B., Kiszler, T., Klocke, D., Kluft, L., Koldunov, N., Kölling, T., Kolstela, J., Kornblueh, L., Kosukhin, S., Lacima-Nadolnik, A., Leal Rojas, J. J., Lehtiranta, J., Lunttila, T., Luoma, A., Manninen, P., Medvedev, A., Milinski, S., Mohammed, A. O. A., Müller, S., Naryanappa, D., Nazarova, N., Niemelä, S., Niraula, B., Nortamo, H., Nummelin, A., Nurisso, M., Ortega, P., Paronuzzi, S., Pedruzo Bagazgoitia, X., Pelletier, C., Peña, C., Polade, S., Pradhan, H., Quintanilla, R., Quintino, T., Rackow, T., Räisänen, J., Rajput, M. M., Redler, R., Reuter, B., Rocha Monteiro, N., Roura-Adserias, F., Ruppert, S., Sayed, S., Schnur, R., Sharma, T., Sidorenko, D., Sievi-Korte, O., Soret, A., Steger, C., Stevens, B., Streffing, J., Sunny, J., Tenorio, L., Thober, S., Tigerstedt, U., Tinto, O., Tonttila, J., Tuomenvirta, H., Tuppi, L., Van Thielen, G., Vitali, E., von Hardenberg, J., Wagner, I., Wedi, N., Wehner, J., Willner, S., Yepes-Arbós, X., Ziemen, F., and Zimmermann, J.: The Destination Earth digital twin for climate change adaptation, EGUsphere, 2025, 1–41, https://doi.org/10.5194/egusphere-2025-2198, 2025.

Gunther, N. J.: A General Theory of Computational Scalability Based on Rational Functions, arXiv, https://doi.org/10.48550/arxiv.0808.1431, 2008.

Hanke, M., Redler, R., Holfeld, T., and Yastremsky, M.: YAC 1.2.0: new aspects for coupling software in Earth system modelling, Geoscientific Model Development, 9, 2755–2769, https://doi.org/10.5194/gmd-9-2755-2016, 2016.

Huijnen, V., Le Sager, P., Köhler, M. O., Carver, G., Rémy, S., Flemming, J., Chabrillat, S., Errera, Q., and van Noije, T.: OpenIFS/AC: atmospheric chemistry and aerosol in OpenIFS 43r3, Geoscientific Model Development, 15, 6221–6241, https://doi.org/10.5194/gmd-15-6221-2022, 2022.

https://doi.org/10.5194/egusphere-2026-1661
Preprint. Discussion started: 17 June 2026

Irrmann, G., Masson, S., Maisonnave, E., Guibert, D., and Raffin, E.: Improving ocean modeling software NEMO 4.0 benchmarking and communication efficiency, Geoscientific Model Development, 15, 1567–1582, https://doi.org/10.5194/gmd-15-1567-2022, 2022.

Madec, G., Bell, M., Benshila, R., Blaker, A., Boudrallé-Badie, R., Bricaud, C., Bruciaferri, D., Carneiro, D., Castrillo, M., Calvert, D., Chanut, J., Clementi, E., Coward, A., de Lavergne, C., Dobricic, S., Epicoco, I., Éthé, C., Fiedler, E., Ford, D., Furner, R., Ganderton, J., Graham, T., Harle, J., Hutchinson, K., Iovino, D., King, R., Lea, D., Levy, C., Lovato, T., Maisonnave, E., Mak, J., Sanchez, J. M. C., Martin, M., Martin, N., Martins, D., Masson, S., Mathiot, P., Mele, F., Mocavero, S., Moulin, A., Müller, S., Nurser, G., Oddo, P., Paronuzzi, S., Paul, J., Peltier, M., Person, R., Rousset, C., Rynders, S., Samson, G., Schroeder, D., Storkey, D., Storto, A., Téchené, S., Vancoppenolle, M., and Wilson, C.: NEMO Ocean Engine Reference Manual, https://doi.org/10.5281/zenodo.14515373, 2024.

Maisonnave, E. and Bourdallé-Badie, R.: Coupling NEMO global ocean with hemispheric Arctic and Antarctic ice models, Tech. Rep. TR/CMGC/22/18, CERFACS, Toulouse, France, https://cerfacs.fr/wp-content/uploads/2022/03/TR-CMGC-22-18.pdf, discusses advanced SAS applications and hemispheric parallelization, 2022.

Maisonnave, E. and Masson, S.: Ocean/sea-ice macro task parallelism in NEMO, Tech. Rep. TR-CMGC-15-54, CERFACS, Toulouse, France, https://cerfacs.fr/wp-content/uploads/2016/03/glo-tr-cmgc-15-54.pdf, technical Report focusing on the OPA-SAS decoupling using OASIS, 2015.

Mechoso, C. R., An, S.-I., and Valcke, S.: Atmosphere-ocean Modeling: Coupling and Couplers, World Scientific, 2021.

Mogensen, K., Keeley, S., and Towers, P.: Coupling of the NEMO and IFS models in a single executable, ECMWF Technical Memorandum 673, European Centre for Medium-Range Weather Forecasts, https://doi.org/10.21957/rfplwzuol, 2012.

Palmer, T. N.: The physics of numerical analysis: a climate modelling case study, Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences, 378, 20190 058, https://doi.org/10.1098/rsta.2019.0058, 2020.

Remy, S. and Bock, J.: Improvement of the SO2 to SO4 conversion processes in IFS-AER and IFS-GLOMAP, Tech. rep., Copernicus Atmosphere Monitoring Service, https://doi.org/10.24380/F84M-3X40, 2018.

Segura, H., Pedruzo-Bagazgoitia, X., Weiss, P., Müller, S. K., Rackow, T., Lee, J., et al.: nextGEMS: entering the era of kilometer-scale Earth system modeling, EGUsphere, https://doi.org/10.5194/egusphere-2025-509, 2025.

Sparrow, S., Bowery, A., Carver, G. D., Christensen, H. M., Hirooka, S., Lott, F. C., Matsunuma, T., Miller, J., Milner, H., Mizuta, R., Murgatroyd, A., O'Reilly, C., Petch, J., Schiemann, R., Wallom, D., Watson, P. A. G., Weisheimer, A., Williams, J., Yamazaki, K., Yang, G.-Y., and Allen, M. R.: OpenIFS@home version 1: a citizen science project for ensemble weather and climate forecasting, Geoscientific Model Development, 14, 3473–3486, https://doi.org/10.5194/gmd-14-3473-2021, 2021.

Tegen, I., Neubauer, D., Ferrachat, S., Siegenthaler-Le Drian, C., Bey, I., Schutgens, N., Stier, P., Watson-Parris, D., Stanelle, T., Schmidt, H., Rast, S., Kokkola, H., Schultz, M., Schroeder, S., Daskalakis, N., Barthel, S., Heinold, B., and Lohmann, U.: The global aerosol–climate model ECHAM6.3–HAM2.3 – Part 1: Aerosol evaluation, Geoscientific Model Development, 12, 1643–1677, https://doi.org/10.5194/gmd-12-1643-2019, 2019.

Tintó Prims, O. and Palomas, S.: Escalador, https://doi.org/10.5281/zenodo.19187009, 2026.

Tintó Prims, O., Castrillo, M., Acosta, M. C., Mula-Valls, O., Sanchez Lorente, A., Serradell, K., Cortés, A., and Doblas-Reyes, F. J.: Finding, analysing and solving MPI communication bottlenecks in Earth System models, Journal of Computational Science, 36, 100 864, https://doi.org/https://doi.org/10.1016/j.jocs.2018.04.015, 2019.

Tintó Prims, O., Palomas, S., Vacondio, S., and Acosta, M. C.: Reproduction scripts for: Quantifying the scaling penalty of sequential coupling, https://doi.org/10.5281/zenodo.19187264, 2026.

https://doi.org/10.5194/egusphere-2026-1661
Preprint. Discussion started: 17 June 2026

Van Noije, T., Bergman, T., Le Sager, P., O'Donnell, D., Makkonen, R., Gonçalves-Ageitos, M., Döscher, R., Fladrich, U., Von Hardenberg, J., Keskinen, J.-P., Korhonen, H., Laakso, A., Myriokefalitakis, S., Ollinaho, P., Pérez García-Pando, C., Reerink, T., Schrödner, R., Wyser, K., and Yang, S.: EC-Earth3-AerChem: a global climate model with interactive aerosols and atmospheric chemistry participating in CMIP6, Geoscientific Model Development, 14, 5637–5668, https://doi.org/10.5194/gmd-14-5637-2021, 2021.

van Noije, T. P. C., Le Sager, P., Segers, A. J., van Velthoven, P. F. J., Krol, M. C., Hazeleger, W., Williams, A. G., and Chambers, S. D.: Simulation of tropospheric chemistry and aerosols with the climate model EC-Earth, Geoscientific Model Development, 7, 2435–2475, https://doi.org/10.5194/gmd-7-2435-2014, 2014.

Vancoppenolle, M., Rousset, C., Blockley, E., Aksenov, Y., Feltham, D., Fichefet, T., Garric, G., Guémas, V., Iovino, D., Keeley, S., Madec, G., Massonnet, F., Ridley, J., Schroeder, D., and Tietsche, S.: SI3, the NEMO Sea Ice Engine, https://doi.org/10.5281/zenodo.7534900, 2023.

Vignati, E., Wilson, J., and Stier, P.: M7: An efficient size-resolved aerosol microphysics module for large-scale aerosol transport models, Journal of Geophysical Research: Atmospheres, 109, https://doi.org/10.1029/2003JD004485, _eprint: https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1029/2003JD004485, 2004.

Wedi, N., Bauer, P., Sandu, I., Hoffmann, J., Sheridan, S., Cereceda, R., Quintino, T., Thiemert, D., and Geenen, T.: Destination Earth: High-Performance Computing for Weather and Climate , Computing in Science & Engineering, 24, 29–37, https://doi.org/10.1109/MCSE.2023.3260519, 2022.

Code availability. The core library (Escalador) is available under the Apache License 2.0 and is archived on Zenodo at https://doi.org/10.5281/zenodo.19187009 (Tintó Prims and Palomas, 2026). The specific scripts and configuration files required to reproduce the analyses and figures presented in this manuscript are also available on Zenodo at https://doi.org/10.5281/zenodo.19187264 (Tintó Prims et al., 2026).

Author contributions. **OT**: Conceptualization, Methodology, Software, Investigation (Use cases 1 and 2), Writing – original draft. **SP**: Conceptualization, Investigation (Use case 3), Writing – review & editing. **SV**: Investigation (Use case 3), Writing – review & editing. **MA**: Supervision, Writing – review & editing.

Competing interests. The authors declare that they have no conflict of interest.

Acknowledgements. Oriol Tintó Prims acknowledges his AI4S fellowship within the "Generación D" initiative by Red.es, Ministerio para la Transformación Digital y de la Función Pública, for talent attraction (C005/24-ED CV1), funded by NextGenerationEU through PRTR.

Researchers Mario Acosta and Sergi Palomas were supported by the European High Performance Computing Joint Undertaking (EuroHPC JU) and the European Union (EU) through the ESiWACE3 project under grant agreement No. 101093054.

Simone Vacondio was supported by Destination Earth.

https://doi.org/10.5194/egusphere-2026-1661
Preprint. Discussion started: 17 June 2026

We acknowledge fruitful discussions with Maria Gonçalves, Philippe Le Sager, and Twan van Noije for providing insight into the OIFS–M7 configuration.