The location of the thermodynamic atmosphere-ice interface in fully-coupled models – a case study using JULES and CICE

Alex E. West[1], Alison J. McLaren[1], Helene T. Hewitt[1], Martin J. Best[1]

[1]Met Office Hadley Centre, Exeter Devon

Correspondence to: Alex E. West (alex.west@metoffice.gov.uk)

Abstract. In fully-coupled climate models, it is now normal to include a sea ice component with multiple layers, each having their own temperature. When coupling this component to an atmosphere model, it is more common for surface variables to be calculated in the sea ice component of the model, the equivalent of placing an interface immediately above the surface. This study uses a one-dimensional (1D) version of the Los Alamos sea ice model (CICE) thermodynamic solver and the Met Office atmospheric surface exchange solver (JULES) to compare this method with that of allowing the surface variables to be calculated instead in the atmosphere, the equivalent of placing an interface immediately below the surface.

The model is forced with a sensible heat flux derived from a sinusoidally varying near-surface air temperature. The two coupling methods are tested first with a 1-h coupling frequency, and then a 3-h coupling frequency, both commonly-used. With an above-surface interface, the resulting surface temperature and flux cycles contain large phase and amplitude errors, as well as having a very 'blocky' shape. The simulation of both quantities is greatly improved when the interface is instead placed within the top ice layer, allowing surface variables to be calculated on the shorter timescale of the atmosphere. There is also an unexpected slight improvement in the simulation of the top-layer ice temperature by the ice model. The surface flux improvement remains when a snow layer is added to the ice, and when the wind speed is increased. The study concludes with a discussion of the implications of these results to three-dimensional modelling. An appendix examines the stability of the alternative method of coupling under various physically realistic scenarios.

1 Introduction

Sea ice has long been recognised as an important component of the climate system, and all climate models taking part in the CMIP5 project now include a sea ice component. Much progress has been made in sea ice modelling since the 1970s. Maykut and Untersteiner (1971) derived governing equations of sea ice thermodynamics, with temperature and salinity-dependent heat capacity and conductivity, and allowing for a snow layer above the ice. Semtner (1975) devised a simple numerical model of sea ice thermodynamics based on a simplification of the Maykut and Untersteiner equations, designed for incorporation in coupled climate models. An appendix to Semtner's study detailed an even simpler model in which the ice had no heat capacity at all, the so-called 'zero-layer' method. The simulation of the spatial coverage of sea ice by even this highly simplified model was found to be reasonably accurate; for example, Johns et al. (2006) and Gordon et al. (2000)

describe the sea ice simulations of HadGEM1 and HadCM3 respectively, both coupled models incorporating this scheme. Hence this method became the basis of the thermodynamics of many sea ice models, with its low computational costs.

As computing power increases, however, the multi-layer model of Semtner is becoming the more commonly-used version. In particular, the Los Alamos sea ice model CICE (Hunke et al., 2013), which is the focus of the present study, bases its thermodynamics on a more complex multi-layer discretisation of the Maykut and Untersteiner equations, as updated by Bitz and Lipscomb (1999), with heat capacity and conductivity fully dependent on salinity and temperature.

Currently, the configuration of models used for climate projections at the Met Office Hadley Centre (HadGEM3) uses the zero-layer version of CICE (Hewitt et al, 2011). The present study arose out of a desire to couple the multi-layer version of CICE to the Met Office atmosphere model, the Unified Model (UM) (Walters et al., 2011), and in particular its surface exchange scheme JULES (Best et al., 2011). Both CICE and JULES perform integrations using a forwards-implicit timestepping method, with much greater stability than would be associated with an explicit calculation; in CICE new ice temperatures are calculated, based on future values of temperature, conductivity, and heat capacity, while in JULES surface temperature and fluxes are calculated, based on future values of the surface exchange coefficients. CICE calculates temperatures for each of the individual ice layers, and the ice surface; JULES calculates all surface variables. Hence a conflict arises when trying to couple the two components; each 'wants' to calculate the surface variables itself, but in practice only one must be allowed to do so, as two different values of surface variables would be associated with two subsequently different model evolutions.

At the root of the problem is that whereas in physical reality the ice and atmospheric temperatures are intimately related, and vary in one system, in the model an explicit interface must be placed between them. Ideally one would solve implicitly for the whole ice and atmosphere column, but in practice while the two systems are separately implicit, the coupling across the interface must be explicit. CICE assumes the interface to lie above the ice surface; the JULES surface exchange scheme assumes it to lie below the ice surface. (Note that the same problem does not arise for the ice and ocean systems, because the base of the ice is at present always assumed to be at the freezing point of seawater).

One possible solution is to place the entire ice column within the atmospheric thermodynamic solver, equivalent to locating the thermodynamic interface at the base of the ice, but this approach would necessitate passing a very large number of fields between the two models, and has been deemed impractical at the Met Office. It is also, in theory, possible to design an implicit scheme for atmosphere and ice in which the two thermodynamic solvers are in different code bases, but in this case it is necessary to pass information between the two components at an instant while the solvers are calculating new temperatures. This would require coupling every atmospheric timestep, which would be too computationally expensive.

The purpose of this study is to examine the two coupling methods under idealised conditions, using a one-dimensional version of the CICE temperature solver, and a miniature version of the JULES surface exchange scheme, under realistic timestep lengths, coupling period lengths, and vertical resolutions, and in particular to determine which gives the more accurate simulation. In section 2, the CICE thermodynamic solver and the JULES surface energy balance solver are described in more detail, along with the two coupling methods. In section 3, the performance of the CICE temperature

solver is examined using its own coupling method, under varying vertical and temporal resolutions. In section 4, the CICE and JULES components are run together using the two different coupling methods, under a variety of different conditions, and the results compared. Finally, in section 5 we discuss the results, and their applicability to fully-coupled models. In the Appendix, the stability of the alternative coupling method under the limits of physically realistic conditions is examined.

2 Description of the models and experiments

2.1 The models: CICE

The fundamental equation solved by the CICE temperature solver is the heat diffusion equation:

$$\rho c_p(S,T)\frac{\partial T}{\partial t} = \frac{\partial}{\partial z}\left(k(S,T)\frac{\partial T}{\partial z}\right) \quad (1)$$

where ρ, cp, S, T, t, z, and k denote ice density, heat capacity, salinity, temperature, time, depth and conductivity respectively. CICE includes an additional term representing penetrating solar radiation, which we neglect for the purposes of this study. Conductivity and heat capacity are parametrised as

$$k(S,T) = K_0 + \frac{\beta S}{T} \quad (2)$$

where

$$K_0 = 2.03 W m^{-1} K^{-1} \quad (3)$$

and

$$\beta = 0.13 W m^{-1} \quad (4)$$

after Untersteiner (1964) and

$$c_p = c_{p0} + \frac{L_0 \mu S}{T^2} \quad (5)$$

where

$$c_{p0} = 2106 J kg^{-1} K^{-1}, \quad (6)$$

$$L_0 = 3.354 \times 10 \ J kg^{-1} \quad (7)$$

and

$$\mu = 0.054 K \ . \quad (8)$$

after Ono (1967), respectively.

The heat diffusion equation is discretised by splitting the ice into N layers of thickness h_i, and using finite timestepping in the usual way. To ensure stability, temperatures are updated using variables from the next timestep, the so-called 'implicit' method:

$$\rho_k c_{pk} \frac{T_k^{m+1} - T_k^m}{\Delta t} = \frac{1}{h_k}\left[K_k\left(T_{k-1}^{m+1} - T_k^{m+1}\right) - K_{k+1}\left(T_k^{m+1} - T_{k+1}^{m+1}\right)\right] \quad (9)$$

where the subscripts m and k denote timestep number and vertical layer number respectively, and $K_k = \frac{2k_{k-1}k_k}{k_{k-1}h_k + k_k h_{k-1}}$ is the 'effective conductivity' at the interface between layers k and $k-1$.

There is an additional equation for the change in surface temperature, T_{sf}:

$$F_0^* + \left(\frac{dF_0}{dT_{sf}}\right) \bullet \left(T_{sf}^{m+1} - T_{sf}^*\right) = K_1\left(T_{sf}^{m+1} - T_1^{m+1}\right) \quad (10)$$

Here, F_0^* represents the sum of radiative, sensible and latent heat fluxes arriving at the ice surface from above; in the absence of melting this is equal to *Fcondtop*, the conductive flux travelling downwards into the ice.

In this way a linear system of equations for the new layer temperatures (plus the surface temperature) is created, $\mathbf{A}\boldsymbol{T}_{new} = \boldsymbol{R}$, where $\mathbf{A}$ is a tridiagonal matrix and $\boldsymbol{T}_{\mathbf{new}}$ is the vector of new layer temperatures. The parameter c_{pk} depends itself upon the layer temperature, T_k^m; in addition, Equation (10) is an approximation, as in reality upwelling longwave radiation has a nonlinear dependence on surface temperature. Because of these factors, it is necessary to iterate the linear solver, updating outgoing longwave radiation, F and c_{pk} at each iteration, to achieve an accurate and energy-conserving solution. (Note that although k also depends upon T, as this variable carries no direct implications for energy conservation it is not updated at each iteration.) CICE allows up to 100 iterations, although generally fewer than 10 will suffice to reduce the energy imbalance to acceptable levels. Hence, in equation 10, starred variables represent variables from the preceding iteration.

It should be noted that CICE also allows for the presence of a snow layer on top of the ice, which introduces an extra row into the matrix equation, with accordingly different heat capacity and conductivity. For this study, however, we assume initially that no snow is present.

2.2 The models: JULES

The principal function of the surface-exchange scheme JULES is to solve the surface energy balance equation, in which a surface temperature is calculated such that incoming fluxes of shortwave and longwave radiation are in balance with outgoing turbulent, radiative and conductive fluxes:

$$(1-\alpha)SW_{in} + LW_{in} - \varepsilon\sigma T_{sfc}^4 + F_{sens}\left(T_{sfc}, T_{air}\right) + F_{lat}\left(T_{sfc}, T_{air}, q_{air}\right) = \kappa_{eff}\left(T_{sfc} - T_{ice}\right) + F_{melt} \quad (11)$$

In this equation SW_{in}, LW_{in} refer to the incoming shortwave and longwave fluxes respectively; F_{sens} and F_{lat} to the net inward sensible and latent heat fluxes respectively; T_{sfc}, T_{air} and T_{ice} to surface temperature, lowest-layer air temperature and uppermost layer ice temperature respectively, q_{air} to lowest layer air specific humidity, k_{eff} to effective conductivity of the top ice layer, α to surface albedo and F_{melt} to the sea ice melt flux. JULES solves this equation by first calculating a 'first guess' explicit solution, calculating fluxes and surface temperature based on surface temperature at the previous timestep, and then calculating an implicit updated solution, in which the exchange coefficients are modified by considering the initial solution. JULES computes surface exchange coefficients over sea ice using the same method as is used over land, as described in Section 2.1 of Best et al (2011). Because the surface temperature simulation carries no implications for energy conservation, the calculation is not iterated.

2.3 The coupling methods and experiments

In their standard formulations, both the CICE thermodynamic solver and the JULES surface exchange solver calculate surface variables. The two coupling methods under investigation arise from opposite methods of resolving this redundancy.

In the standard 'CICE' coupling method (Figure 1a), the atmosphere, or surface exchange scheme, calculates fluxes of incoming shortwave and longwave radiation based on the evolving atmospheric state whose lower boundary condition is the ice surface temperature from the previous coupling instant. The atmosphere then averages these over the coupling period, and passes them to CICE at the end of that period. CICE then uses these incoming fluxes throughout the same coupling period in the first row of the tridiagonal matrix equation the row concerning the surface temperature (equation 10), each time iterating the solver until convergence is achieved. In the process, CICE computes the remaining surface fluxes (outgoing radiative, turbulent and conductive fluxes) and hence the net surface flux. This approach is equivalent to placing an interface between JULES and CICE immediately above the ice surface.

In the alternative, 'JULES' coupling method under investigation (Figure 1b), the surface temperature is a prognostic variable of the atmosphere or surface exchange model, and is not passed from CICE; instead, the temperature and effective conductivity (the latter defined as $\frac{2K_1}{h_1}$) of the top ice layer are passed at each coupling instant. The surface exchange scheme calculates an updated surface temperature, along with radiative fluxes, turbulent fluxes, surface ice melt, and downward conductive flux into the top layer of ice from the surface, in a fully implicit boundary layer solution, given these lower boundary conditions. The downward conductive flux and ice melt flux are averaged over a coupling period and passed to CICE for use in the same coupling period. CICE proceeds to solve the tridiagonal matrix occasion in the normal way, except that the top row of the equation is removed; the downwards conductive flux provided by the surface exchange scheme is then used as forcing for the top ice layer. At the end of the coupling period, the new temperature and effective conductivity of the top ice layer are passed back to the atmosphere. This approach is equivalent to placing an interface between JULES and CICE immediately below the ice surface.

It should also be noted that in HadGEM3, fluxes are always passed as gridbox means, to ensure conservation. This point only becomes relevant in 3D modelling, where sea ice may cover only part of a grid cell; in this case, the relevant flux is multiplied by the grid cell ice concentration before being passed to the coupler for regridding. This is necessary because of the parallel coupling of HadGEM3 (see Section 4.4); underlying ice concentration may change during a coupling period, and hence the amount of energy being passed must be correctly represented by multiplying, effectively, by the area over which it is valid.

3 Testing the impact of varying resolution on an idealised solver

3.1 Setup

In this experiment, the penetrating solar radiation term was ignored, and the ice was assumed to be fresh, in order that the conductivity and specific heat capacity are constant. The ice was assumed to be 1m thick, and there is no snow cover. The diffusion equation was forced at the top of the ice by a sinusoidally varying heat flux:

$$F_{sf} = A\,\mathrm{Re}\left(\exp i\omega t\right) \tag{12}$$

There exists an exact analytical solution to the diffusion equation with this surface forcing, for an infinitely deep ice cover (after Best et al, 2005):

$$T = T_B + \frac{A}{\sqrt{\rho c \omega \kappa}}\,\mathrm{Re}\left(\exp\frac{-z}{l}\exp i\left(\frac{-z}{l} + \omega t - \frac{\pi}{4}\right)\right) \tag{13}$$

Where $T \to T_B$ as $z \to \infty$. $l = \sqrt{\frac{2\kappa}{\rho c \omega}}$ is the e-folding depth.

This analytical solution was compared to the solution from the CICE temperature solver under 6 different conditions, summarised in Table 1. In these experiments, the timestep length, coupling period length and vertical resolution were varied, from extremely low values designed to give results as close as possible to 'truth', to higher values considered to be typical of coupled model experiments.

3.2 Results

Figure 2 displays the simulation of two key variables by the temperature solver: the surface temperature, and the temperature at a depth of .125m (roughly analogous to the top layer temperature in standard CICE, which uses 4 vertical layers). It is clear that under very high temporal and vertical resolution, CICE produces a simulation that is virtually indistinguishable from the analytic solution. As one would expect, when these resolutions are reduced to more realistic levels inaccuracies appear.

When the timestep length is increased to 1 hour (but the high vertical resolution is maintained), there is a slight increase in the error of the surface temperature simulation, which is still very small in proportion to the cycle amplitude. For the .125m temperature, a small phase lead of around 30 minutes is introduced, and the amplitude is reduced by a tiny amount (0.02°C); the diurnal cycle of .125m temperature error has an amplitude of about 0.03°C.

The effect of decreasing the vertical resolution is more marked. For the surface temperature, we see a a large phase lead introduced, of 90 minutes, but also a marked increase in amplitude, from 1.2 to 1.5◦ C; this results in some comparatively high errors, of up to 0.6°C. On the other hand, the diurnal cycle of .125m temperature is reduced in amplitude slightly, and has a lower phase lead of about 1 hour. The errors have magnitude of up to 0.09°C. The contrasting effects of the decreased vertical resolution on surface and top layer temperature can be understood by considering that the surface temperature is forced by the air temperature, and damped by the ice temperature. A top ice layer of thickness 1cm can warm or cool more easily for a given forcing than can a top ice layer of thickness 25cm; therefore, when the entire top 25cm of ice has to vary in unison, the amplitude of its cycle is reduced, and its damping effect on the surface temperature is correspondingly reduced, which can hence vary more strongly in response to the air temperature forcing.

Lastly, we look at the effects of moving to a 3-hour coupling period, with timestep length maintained at 1 hour (Figure 3). It is apparent that this change has little effect on the phase or amplitude of the surface temperature simulation, and only serves to make the diurnal cycle more 'jagged'; at each coupling period, indicated by the vertical solid grey lines, the surface temperature jumps by a large amount, and over the following two (non-coupling) timesteps, moves backwards by a smaller amount as the sea ice adjusts towards a new equilibrium.

The error in the 4-layer experiment should give cause for concern, as this is a fairly realistic resolution for most implementations of CICE in coupled models. In the next section, therefore, we compare the simulations at realistic resolution, using the two different coupling methods.

4 Comparing the two coupling methods under realistic resolution

4.1 Setup

For this experiment, the solver was run under 6 different setups. Firstly, two 'control' experiments were undertaken, in which the ice, atmosphere and coupling timesteps were each 1 second. In the first control, the ice was given 100 layers, to provide a 'truth' against which to compare subsequent experiments; in the second control, the ice was given 4 layers, to separate the effects of high timestep values from the effects of low vertical resolution. The two control experiments were run using the 'CICE' coupling method, with the surface variables calculated by the ice model, but at these timestep values the coupling method has negligible impact on the simulation.

The solver was then run with 4 vertical layers, an ice timestep of 1 hour, atmosphere timestep 20 minutes, and coupling period of 1 hour, fairly realistic values for a coupled model, using the two different coupling methods, 'CICE' and 'JULES'.

A further two experiments were then performed, using a coupling period of 3 hours, also a common period found in coupled model runs.

The solver was forced with incoming sensible heat flux only, driven by a diurnal cycle of atmospheric surface temperature $T_{atmos} = A_T \exp i\omega t$, with wind speed set to 5 m/s. For the 'CICE' coupling, T_{atmos} is averaged over a coupling period and passed to the ice model, which calculates from this the incoming sensible heat flux, and uses this as forcing for the temperature solver to calculate internal and surface ice temperatures. For the 'JULES' coupling, a self-contained 'atmosphere model' uses T_{atmos} and T_1 (top-layer ice temperature) to implicitly calculate surface fluxes, including $F_{condtop}$, downwards conductive flux, accumulates and averages this over the coupling period and passes it to the ice model as forcing for the solver.

4.2 Results

4.2.1 1-h coupling

Figure 4 displays the simulation of key variables by the high-resolution control runs and by the test runs, using a 1-hour coupling period length. The forcing atmospheric temperature is indicated in Figure 4a. First examining the surface flux (Figure 4b), we compare the two control runs and note that the decrease in vertical resolution is associated with a slight decrease in amplitude and a phase lag. We then see that when the 'JULES' coupling method is used, there is little further error associated with the decrease in temporal resolution (blue line). When the 'CICE' coupling method is used, however, there is an additional phase lag and amplitude decrease, and in addition the diurnal cycle becomes more jagged.

Interpreting these results, it is likely that the additional phase lag is a consequence of the atmosphere model 'seeing' a surface flux calculated in the previous CICE coupling period, which is itself based on an atmospheric temperature valid for the period before that, up to 2 hours previously. With the 'JULES' method, however, the surface flux is able to respond immediately to the changing atmospheric temperature. There is a corresponding delay in the atmosphere model 'sensing' the damping response of the top layer ice temperature to the changing surface flux. However, the resulting phase lead is tiny in comparison to the phase lag of the 'CICE' method.

We now consider the atmosphere model surface temperature (Figure 4c). In this variable, a decreasing vertical resolution is associated with an increase in amplitude and a phase lead. Again, using the 'JULES' method, a decreasing temporal resolution makes little difference, causing a tiny phase lag and a slightly less smooth shape compared to the 4-layer control. Using the 'CICE' method produces a much more blocky shape, and a substantial phase lag. However, as the 4-layer control itself has a phase lead relative to the high-resolution control, the 'CICE' method actually has a more accurate phase; the temporal and vertical errors 'cancel', while the 'JULES' method maintains a phase lead.

How the ice model 'sees' the surface temperature is demonstrated in Figure 4d. The diurnal cycle is very similar to that of the atmosphere model surface temperature for the two control runs, due to their low timestep length. The ice model does not have knowledge of the surface temperature in the 'JULES' coupling method and this line is not plotted. The surface

temperature simulation in the 'CICE' method is very similar to the control; the phase lag experienced by the atmosphere model is due to the coupling delay only.

Conversely, Figure 4e demonstrates how the atmosphere model 'sees' the top layer temperature, in the 4-layer control and in the 'JULES' coupling method (as in the 'CICE' method the atmosphere has no knowledge of this variable). There is a slight phase lag relative to the control, and associated jaggedness of the diurnal cycle, owing to the need to hold the temperature constant over each coupling period, rather than update it every atmospheric timestep.

The lower panels (Figure 4f and 4g) compare the internal ice temperatures at .125m (top layer) and .625m (third layer) depth in the four experiments. For both variables, the decrease in vertical resolution is characterised by a decrease in amplitude and a phase lead which are both more severe in the deeper variable. The decrease in timestep length produces additional amplitude decrease and phase lead which are very similar in the two coupling methods. It is interesting to note that the errors are marginally smaller for the 'JULES' method. This is likely due to the fact that in the 'JULES' method, changes in T_{atmos} can propagate quickly downwards to changes to f_{surf} on the 20-minute atmospheric timestep, the main bottleneck occurring in the coupling, as f_{surf} forces changes in T_1 on the slower 1-hour coupling period. In contrast, in the 'CICE' method, each link in the chain – from T_{air}, to T_{sfc} and $f_{condtop}$, to T_1 – must communicate on a slow 1-hour timestep. In consequence, the 'JULES' method simulation is slightly closer to that of the 4-layer control.

4.2.2 3-h coupling

The results of the experiments when 3-hour coupling is used are shown in Figure 5. For the surface flux (Figure 5b), again, decreasing temporal resolution is identified with a small phase lag and amplitude decrease in the 'JULES' method; the simulation is very similar to that with the 1-hour coupling period, although slightly less smooth. For the 'CICE' method, however, the phase lag and amplitude decrease are greatly magnified; the peak of the diurnal cycle occurs 2-3 hours too late, and the cycle has a very discontinuous shape.

Considering the surface temperature (Figure 5c), the 'JULES' method again produces a simulation with a 3-hour coupling period which is quite similar to that with the 1- hour period, though less smooth. Again the effect of the 'CICE' method is to produce a phase lag. Whereas in the 1-hour coupling period case, however, this phase lag almost exactly cancelled the lead of the increased vertical resolution, in the 3-hour case the lead is much greater, and the absolute phase error of the method is actually greater than that of the 'JULES' method, in an interesting demonstration of the dangers of cancelling errors.

When considering the ice variables (Figure 5f and 5g) there are again few clear differences between the simulations, but again the error is marginally smaller for the 'JULES' method than for the 'CICE' method. Again this is likely related to the 'chain' by which changes propagate from T_{atmos}, via T_{sfc} and f_{surf}, to T_1. The 'JULES' method involves a 'fast' link, on the 20-minute atmospheric timestep, from T_{atmos} to T_{sfc} and f_{surf}, and a 'very slow' link, on the 3-hour coupling timestep, from f_{surf} to T_1. By contrast, the 'CICE' method involves a 'very slow' link, on the 3-hour coupling timestep, from T_{atmos} to T_{sfc} and f_{surf}, and a 'slow' link, on the 1-hour CICE timestep, from T_{sfc} to T_1. While the rate of propagation is for both methods

dominated by the 3-hour coupling 'bottleneck', therefore, changes in T_{atmos} are still able to propagate slightly more quickly with the 'JULES' method.

In summary, the deterioration in simulation of the atmospheric variables that is associated with decreased temporal resolution is significantly reduced by using the 'JULES' coupling method. There is also a small improvement in the simulation of the ice variables, although this is very marginal.

4.3 Varying the parameters of the experiments

To gain some idea of the generality of the results, the parameters of the experiment were varied. Firstly, the coupling methods were tested with an 11cm snow layer present above the ice. Secondly, they were tested without a snow layer, but with the wind speed increased from 5 m/s to 20 m/s, to examine the impact of strengthening the coupling between the forcing air temperature and the surface.

The results are presented in Figure 6a-6f, in which for each additional experiment, and for the original experiment, the surface flux and the top layer temperature are plotted. (For the snow layer experiment, the top layer temperature corresponds to the snow layer temperature; for the other experiments, to the top ice layer temperature). For clarity, only the two 'control' experiments and the two 3-h coupling experiments are shown.

Looking first at the 'snow layer' experiment, it can be seen that the surface flux diurnal cycle displays greatly reduced amplitude in all setups (Figure 6a, 6c), a consequence of the extra insulation provided by the snow layer decreasing conduction through the ice. The 'JULES' method here displays a slight amplitude increase relative to the 4-layer control, a result of the 3 hour delay in the atmosphere 'sensing' the damping snow layer temperature, and thus allowing the surface flux to overshoot. However, the errors are still far smaller than those of the 'CICE' method. The snow layer temperature (Figure 6b, 6d) has a much larger diurnal cycle than does the top ice layer temperature in the original experiment, due to its much lower heat capacity. Relative to the 4-layer control, the 'JULES' method overestimates the amplitude, while the 'CICE' method underestimates the amplitude, precisely as is the case for the surface flux.

Next examining the 'wind speed' experiment, in this case the surface flux diurnal cycle (Figure 6e) is greatly increased in magnitude under all setups, as the increased wind speed facilitates heat loss from the surface to the air (and vice versa). Similarly to the snow layer experiment, the 'JULES' method develops an anomalously high amplitude, related to the persistent 'overshoot' in surface flux during each coupling period. This is because the rate at which the surface flux changes during each coupling period is directly proportional to the wind speed, and is therefore four times greater in this experiment; the overall surface flux amplitude, although larger than in the 5 m/s experiment, does not increase in direct proportion. Hence the overshoot is higher in proportion here. However, the 'JULES' method errors are still considerably lower than those of the 'CICE' method. For the top layer ice temperature (Figure 6f), both methods produce simulations very close to those of the 4-layer control.

4.4 Serial versus parallel coupling

In the experiments described above, the forcing flux being passed from the atmosphere to the ice was used as forcing for the ice model for the same coupling period as that during which it was calculated. This is a framework in which atmosphere and ice models are run in sequence, the so-called 'serial' coupling method. While many coupled models function in this way, some (including HadGEM3) use the alternative 'parellel' method, in which atmosphere and ice models are run concurrently. This entails that the atmosphere-to-ice forcing flux is used instead as forcing for the ice during the coupling period after that in which it was calculated. The serial and parallel frameworks are demonstrated schematically in Figure 7.

The parallel method can be more computationally efficient, but is less accurate (as is demonstrated below). The tests of Section 4.1-4.3 were carried out below using the series method, despite HadGEM3 using the parallel method, in order to eliminate the additional source of inaccuracy caused by parallel coupling and therefore enable the results to be more relevant to the wider community. However, it is also useful to compare the relative performance of the 'CICE' and 'JULES' methods under parallel conditions. The tests were therefore carried out again, with the 1D solver edited to mimic a parallel system, rather than a serial one.

The results are shown in Figure 6a-6b (serial coupling) and Figure 6g-6h (parallel coupling) As seen in the sensitivity experiments of Section 4.3, the 'JULES' method displays a deterioration in the surface flux simulation relative to the control (series coupling), shown in Figure 6g, with the surface flux 'overshoot' again enhanced, and the amplitude increased accordingly, a result of the extra 3 hours of delay in the atmosphere receiving a damping response from the top layer temperature to the original overshoot. The reason that the 'CICE' method does not display a similar deterioration is that for this method, one of the variables immediately adjacent to the interface (the air temperature), is not free to vary in this experiment, but prescribed, and therefore the atmosphere 'pays no penalty' for the delay in receiving the forcing from below – a situation which would not occur in a fully coupled model. The three hour lag introduced to the top layer temperature simulation (Figure 6h) is noticeable in both the 'JULES' and 'CICE' methods, and incidentally demonstrates, independently of this study, the drawbacks of using parallel coupling as opposed to serial. It should be pointed out however that this drawback is much reduced when 1-h coupling is used; also, that despite the deterioration, the surface flux errors for the 'JULES' method are still substantially lower than those of the 'CICE' method.

5 Discussion and conclusions

This study has compared, under idealised conditions, the performance of the CICE temperature solver under varying resolutions, and using two different methods of coupling with an atmospheric model. It has been shown that low vertical resolution within the ice can be the source of significant errors in simulating the diurnal cycle. It has been shown that in simulating an idealised diurnal cycle of ice temperatures and surface fluxes, a coupled model in which an atmosphere-ice interface is placed within the ice performs considerably better than one in which an interface is placed at the ice surface, under typical temporal and vertical resolutions; the simulation of surface temperature and surface flux are in general

significantly improved, and the simulation of within-ice variables also improves slightly. It is seen that if a thin snow layer is present, or if the wind speed is increased, the 'JULES' method still simulates the surface flux more accurately, although the margin is reduced.

What is the reason for the improved simulation obtained by simulating surface variables within the atmosphere model, rather than the ice model? The root cause is probably that in the experiments, as is usually the case in coupled models and in the real world, the principal thermodynamic forcing on the surface, and the ice, comes from above rather than below. Air temperature and radiation conditions usually change more rapidly than do properties of the underlying ocean, and of the sub-surface ice. Therefore it is not surprising that an improved simulation is obtained by placing a higher proportion of the ice-surface system within the atmosphere model, from which the forcing comes. With a thin snow cover present, or with increased windspeed, the improvements offered by the 'JULES' method grow slightly less. The reason is likely that both modifications have the effect, for the 'JULES' method, of increasing the magnitude of the surface flux response during each coupling period relative to the surface flux amplitude, thus allowing the overestimation of the amplitude to be worsened. There is no corresponding deterioration for the 'CICE' method, as here the surface flux does not change during the coupling period. However, the 'JULES' method still produces substantially lower surface flux errors. It can be concluded that although the top layer temperature simulation is not systematically better in either method, the 'JULES' method produces a better surface flux simulation under most circumstances.

At first sight, this conclusion appears to disagree with the statement in the introduction to Section 2 of the CICE documentation (Hunke et al, 2013), that 'accuracy may be significantly reduced' by solving for surface temperature in the atmosphere model. However, this statement relates specifically to the hypothetical necessity of artificially reducing effective conductivity to ensure stability in such a situation, rather than the inherent accuracy of the coupling method. In practice, we have found that reducing effective conductivity is not necessary (see Appendix).

This prompts the question: how realistic were the conditions under which the one-dimensional experiments were held, and to what degree would this improvement carry across to the simulation of ice and atmospheric variables in a non-idealised setting? Clearly the best way to answer this question would be to test two coupled models, one using each method. However, the differences between the two setups involve substantial structural changes to all components of the HadGEM3 model, and this option was deemed impractical. Following the results of these experiments, the 'JULES' coupling method is being implemented in the Met Office HadGEM3-GC3 coupled model for use with CICE's capability for multilayer thermodynamics, and when this becomes operational there will be an opportunity to compare the simulation of processes over sea ice to other fully-coupled models which use CICE with the standard 'CICE' coupling method. It is nevertheless possible to use the insight provided by the idealised experiments to gain some idea of the likely effects of the different coupling methods in a 3-D simulation.

The principal effect of the 'CICE' coupling, as opposed to the 'JULES' coupling, is to damp and delay the response of the surface flux (equal in these experiments to the sensible heat flux) to changes in surface air temperature. These changes are applied in the experiments as variations of around 5°C in the course of about 12 hours. Variations in air temperature of this

rate and magnitude are common in the Arctic Ocean, although often they occur in response to changes in cloud cover, or the passage of frontal boundaries, rather than to the diurnal cycle (e.g. Persson et al, 2002). Nevertheless, the implication of the 1-D experiments is that a model using the 'CICE' coupling method will simulate a surface flux response that is overly delayed and damped, relative to a model using the 'JULES' coupling. In effect, the coupling between the atmosphere and the underlying sea ice is weaker, and the atmosphere is likely to behave more like an isolated system.

The effects of this would be complex. A mild airmass moving over cold sea ice tends to be diabatically cooled at the surface via the surface flux response, while the opposite will occur when a cold airmass moves over less cold sea ice. A delayed and damped surface flux response would tend to reduce the rate of modification of airmasses, allowing them to retain characteristics for longer. A similar effect would be likely to be seen in the event of air temperatures responding to changes in radiative forcing due to cloud cover. Normally, the response of surface flux would likely be to moderate diabatic heating or cooling of air masses due to these radiative effects, by transferring some of this heating or cooling into the sea ice; a delayed, damped response would hinder this modification. In this way, it is possible that anomalous characteristics of neighbouring airmasses would become more exaggerated, relative to the real world, when using the 'CICE' coupling method, with unpredictable consequences for atmospheric dynamics. The perturbed parameter experiments demonstrate that under very windy, stormy conditions, the reverse effect might be seen; the surface flux could respond too quickly, and too strongly, thus allowing airmass modification to take place too quickly.

It is seen in Section 4 that the choice of coupling method has little direct impact on the internal sea ice simulation. However, the sea ice simulation will be strongly affected through the atmospheric response described above, whose dynamics will affect advection of warm and cold air over the ice, as well as advection of the ice itself. As the 'JULES' coupling method produces a more realistic surface flux response to changes in air temperature, it appears clear that, all other factors being equal, this coupling method would simulate a more accurate evolution of atmosphere and sea ice.

A secondary finding of this study has been that the vertical resolution at the top of the sea ice is of similar importance to the coupling method used in terms of simulating a realistic surface flux, as demonstrated in Figure 4b. In the current configuration of CICE, whereby all layers are equally spaced within the ice, this implies that surface flux response will tend to be stronger, and more realistic, in regions of thin ice. This suggests that the implementation of variably-spaced layers, with higher resolution near the top of the ice, would be a logical objective to pursue subsequently, to further improve surface flux simulation.

The main focus of this study has been the accuracy of the two coupling methods; a separate question is their stability. The 'CICE' method of coupling is known to have major problems of instability arising from the explicit interface in the surface exchange, an area where processes occur relatively quickly (e.g. Best et al, 2004). However, the 'JULES' method has its own explicit interface, below the ice surface, and is therefore also likely to become unstable under certain conditions. A detailed analysis of the stability of the 'JULES' method in the one-dimensional case is described in Appendix A. The principal factors affecting stability are found to be ice thickness and wind speed; a prediction from this analysis is that setting a minimum ice thickness of 30cm in a coupled model is sufficient to avoid instability in all situations. In practice, however,

in test runs of the coupled model a minimum ice thickness of 20cm has been found to be sufficient to avoid instability. This is probably because in the fully-coupled model, other negative feedbacks are at work in the atmosphere that act to damp oscillations caused by the explicit coupling, and prevent instability.

It is planned to follow this paper with a study examining the simulation of sea ice in HadGEM3 resulting from the implementation of multilayer sea ice, using the 'JULES' coupling method.

Appendix A: Stability of the 'JULES' method of coupling

In this section, the one-dimensional model is used to investigate the conditions under which the solver becomes unstable, prior to its implementation in the Met Office coupled model.

In the stability experiment, the model was run for 5 days; for the first day, the atmospheric temperature was held constant at -20°C, but at the beginning of the second day, the atmospheric temperature was abruptly changed to -15°C; the solver was judged to be stable or unstable according to whether the variables converged to a new solution, and the nature of the convergence was examined. The test was performed under typical modelling conditions of 4 ice layers, ice timestep 1 hour, atmospheric timestep 20 minutes, and of coupling period length 3 hours. The initial parameter that was varied was the ice thickness; the test was performed for six different thicknesses of ice: 1m, 20cm, 10cm, 5cm, 1cm and 1mm. In each case, the top layer ice temperature converged to a new solution, the convergence tending to be most rapid for the thinnest ice (Figure A1).

From this it appears that under normal modelling conditions, the ‚JULES' coupling method is not inherently unstable to sudden perturbations, and tends to be more, rather than less stable, for thin ice. This is perhaps surprising, as it would be thought that thin ice would tend to react more sensitively to perturbations in conductive flux, given its lower thermal inertia. However, counteracting this is the higher effective conductivity of thin ice, meaning that perturbations in top conductive flux will tend to propagate more rapidly through the ice during a coupling period, reducing the resulting change in top layer ice temperature. It also means that as ice thins, the response of the conductive flux comes to dominate the surface energy balance, effectively 'locking' surface temperature to top layer ice temperature, and reducing variation in conductive flux.

To examine the reasons for the stability more carefully, we derive theoretical limits on perturbations to top layer temperature and conductive flux. Given an equilibrium solution to the coupled system $\left(F_{eq}, T_{1\_eq}, T_{sfc\_eq}\right)$, and perturbations around this solution $\left(\hat{F}, \hat{T}_1, \hat{T}_{sfc}\right)$, it can be shown from the surface energy balance equation that the perturbation conductive flux produced by the atmosphere is constrained by the perturbation top layer ice temperature in the following way:

$$\left|\hat{F}\right| \le \frac{k_{eff} OFE}{k_{eff} + OFE} \bullet \left|\hat{T}_1\right| \qquad \text{(A1)}$$

where $k_{eff} = 2k_1 / h_1$ is the effective conductivity of the top layer and $OFE = \dot{f}_{sens}(T_{sfc\_eq}) + \dot{f}_{lat}(T_{sfc\_eq}) + \dot{f}_{rad}(T_{sfc\_eq})$ represents the total rate of change of the surface radiative, sensible and latent heat fluxes with respect to surface temperature at $T_{sfc} = T_{sfc\_eq}$, and tends to reach its highest values under very windy, stormy conditions. It can be seen that the controlling constant here tends to the finite limit *OFE* as $h_1 \to 0$.

Meanwhile, in the ice thermodynamic solver, energy balance considerations provide a constraint on the magnitude of the change in $\hat{T}_1$ during a coupling period:

$$\left|\Delta T_1\right| \le \frac{t_c}{c_p \rho_I h_1}\left|F\right| \tag{A2}$$

where t_c, c_p, ρ_I and h_1 represent coupling period length, ice heat capacity, ice density and top layer ice thickness respectively. This, together with equation (A1), prevents instability for $\frac{t_c OFE}{h_1} < c_p \rho_I$. The system is therefore stable for $h_1 > 5$cm (equivalently total ice thickness < 20cm) in all but the most extreme atmospheric conditions, and for $h_1 > 10$cm (equivalently ice thickness < 40cm) under all realistic atmospheric conditions.

However, for thin ice a second constraint becomes important. A dimensional analysis of the heat diffusion equation for ice shows that with 3 hour coupling, the thermal inertia term can no longer provide the dominant balance to top conductive flux for ice of layer thickness under about 10cm, and becomes negligible for ice of layer thickness under about 2cm, causing the dominant balance in the equation to be between top conductive flux and conduction with the layer below. In this situation, given a top conductive forcing, the ice temperatures will converge very quickly to a linear temperature profile with uniform conductive flux, meaning that

$$\left|\hat{T}_1\right| = \frac{h_1}{2k_1}\left|\hat{F}_1\right| \tag{A3}$$

Combined with equation (A1), this prevents instability completely.

In summary, equations (A1), (A2) and (A3) show that instability cannot occur in the limit of very thick ice (when thermal inertia dominates), due to a highly damped response of top layer temperature to perturbations of conductive flux, and also cannot occur in the limit of very thin ice (when conduction to the ocean dominates), due to the surface temperature becoming virtually 'locked' to top layer ice temperature, perturbations in conductive flux becoming correspondingly small (i.e. when $k_{eff} >> OFE$), and these perturbations very easily propagating through the ice to the ocean. It is noticeable that in Figure A1, the least stable solutions appear to occur for intermediate ice thicknesses (5cm, 10cm), when neither conduction nor thermal inertia dominates, but the 'overlap' in the two conditions is nevertheless sufficient to allow a relatively rapid convergence.

The question arises as to whether the solver would continue to converge for all ice thicknesses were any of the parameters in equations (A1), (A2) or (A3) altered. Parameters c_p, ρ_I and t_c are assumed to be at the lower, lower and upper limits of physical plausibility respectively in equation (2), and to vary them in the opposite direction would serve only to strengthen the limits on convergence. The parameter *OFE*, however, depends strongly on the rate of change of turbulent fluxes with respect to surface temperature, and therefore on wind speed. In the initial stability experiments, wind speed was set to 5 m/s, a fairly typical value for many synoptic situations. Particularly with the passage of extratropical depressions, however, wind speeds can reach much higher values.

The perturbation experiment was repeated, but this time two parameters were varied: ice thickness from 1mm to 1m, and wind speed from 0 m/s to 50 m/s, the upper limit roughly representing the very highest wind speeds possible during extratropical storms. The results are shown in Figure A2. It is seen that the solver is no longer unconditionally stable, with instability setting in at a wind speed of around 23 m/s, at first for a narrow band of ice thicknesses close to 10cm, a band which steadily widens as wind speed increases. At all wind speeds the solver remains stable in the limit of thin ice. However, at the upper limit of wind speed, the solver is unstable for ice thicknesses of between roughly 4cm and 25cm.

This result holds for t_c = 3 hours, but t_c = 1 hour is also a fairly widely used coupling period, and is likely to become more so as computing power increases. The experiment was repeated for t_c = 1 hour (Figure A3). In this case, the solver is stable for all ice thicknesses and wind speeds, although at the upper limit of wind speed, convergence is extremely slow for ice thicknesses of around 7cm. (Clearly the second region of slow convergence, to the right of the figure, is not a concern, as this is caused by higher thermal inertia of thick ice, is entirely physically realistic, and will not lead to instability).

In summary, it is found that the coupled solver system is stable under all physically realistic situations when 1-hour coupling is used, but may become unstable in very windy conditions when 3-hour coupling is used, for certain values of ice thickness.

Author contribution

The 'JULES' method of coupling was originally devised by Martin Best. The one-dimensional experiments of Sections 3 and 4.1 were designed and carried out by Alison McLaren. The perturbation experiments of Section 4.3, the 'parallel coupling' experiments of Section 4.4, and the stability experiments of the Appendix were designed and carried out by Alex West. All results were plotted and analysed by Alex West with advice and assistance from Alison McLaren, Martin Best and Helene Hewitt. The paper was written in its final form by Alex West with input from all three contributing authors.

Code availability

All code used in the production of this study is available as a supplement to the article, together with a file giving instructions for its use (README).

Acknowledgements

The authors thank Ann Keen and Ed Blockley for help and advice prior to submission, and also thank Dirk Notz and Anton Beljaars for their insightful reviews.

This study was supported by the Joint UK DECC/Defra Met Office Hadley Centre Climate Programme (GA01101).

References

Best, M. J., Beljaars, A., Polcher, J., and Viterbo, P.: A Proposed Structure for Coupling Tiled Surfaces with the Planetary Boundary Layer, Journal of Hydrometerorology, 5, 1271–1278, doi: 10.1175/JHM-382.1, 2004

Best, M. J., Cox, P. M., and Warrilow, D. M.: Determining the optimal soil temperature scheme for atmospheric modelling applications, Boundary Layer Meteorology, 114, 111–142, 2005.

Best, M. J., Pryor, M., Clark, D. B., Rooney, G. G., Essery, R. L. H., Menard, C. B., Edwards, J. M., Hendry, M. A., Porson, A., Gedney, N., Mercado, L. M., Sitch, S., Blyth, E., Boucher, O., Cox, P. M., Grimmond, C. S. B., and Harding, R. J.: The Joint UK Land Environment Simulator (JULES), model description - Part 1: Energy and water fluxes, Geoscientific Model Development, 4, 677–699, doi:10.5194/gmd-4-677-2011, 2011.

Bitz, C. M. and Lipscomb, W. H.: An energy-conserving thermodynamic model of sea ice, Journal of Geophysical Research, 104, 15 669–15 677, 1999.

Gordon, C., Cooper, C., Senior, C. A., Banks, H., Gregory, J. M., Johns, T. C., Mitchell, J. F. B. and Wood, R. A., The simulation of SST, sea ice extents and ocean heat transports in a version of the Hadley Centre coupled model without flux adjustments, Climate Dynamics, 16, 147-168, doi:10.1007/s003820050010, 2000.

Hewitt, H. T., Copsey, D., Culverwell, I. D., Harris, C. M., Hill, R. S. R., Keen, A. B., McLaren, A. J., and Hunke, E. C.: Design and implementation of the infrastructure of HadGEM3: the next-generation Met Office climate modelling system, Geosci. Model Dev., 4, 223-253, doi:10.5194/gmd-4-223-2011, 2011.

Hunke, E. C., Lipscomb, W. H., Turner, A. K., Jeffery, N., and Elliott, S.: CICE: the Los Alamos Sea Ice Model, Documentation and Software, Version 5.0, Los Alamos National Laboratory Tech. Rep. LA-CC-06-012, Los Alamos, New Mexico, available at: http://oceans11.lanl.gov/trac/CICE/attachment/wiki/WikiStart/cicedoc.pdf?format=raw (last access: 27 October 2015), 2013. Documentation and Software User's Manual Version 5.0, 2013.

Johns, T. C., Durman, C. F., Banks, H. T., Roberts, M. J., McLaren, A. J., Ridley, J. K., Senior, C. A., Williams, K. D., Jones, A., Rickard, G. J., C. S., Ingram, W. J., Crucifix, M., Sexton, D. M. H., Joshi, M. M., Dong, B.-W., Spencer, H., Hill, R. S. R., Gregory, J. M., Keen, A. B., Pardaens, A. K., Lowe, J. A., Bodas- Salcedo, A., Stark, S., and Searl, Y.: The New

Hadley Centre Climate Model (HadGEM1): Evaluation of Coupled Simulations, Journal of Climate, 19, 1327–1353, doi:http://dx.doi.org/10.1175/JCLI3712.1, 2006.

Maykut, G. A. and Untersteiner, N.: Some Results from a Time-Dependent Thermo- dynamic Model of Sea Ice, Journal of Geophysical Research, 7, 1550–1575, 1971. Ono, N.: Specific heat and fusion of sea ice, Physics of Snow and Ice, 1, 599–610, 1967.

Persson, P. O. G., Fairall, C. W., Andreas, E. L., Guest, P. S., and Perovich, D. K.: Measurements near the Atmospheric Surface Flux Group tower at SHEBA: Near-surface conditions and surface energy budget, Journal of Geophysical Research, 107, doi:10.2019/2000JC000705, 2002.

Semtner, A. J.: A Model for the Thermodynamic Growth of Sea Ice in Numerical Investigations of Climate, Journal of Physical Oceanography, 6, 379–389, 1975.

Untersteiner, N.: Calculations of temperature regime and heat budget of sea ice in the Central Arctic, Journal of Geophysical Research, 69, 755–4766, 1964.

Walters, D. N., Best, M. J., Bushell, A. C., Copsey, D., Edwards, J. M., Falloon, P. D., Harris, C. M., Lock, A. P., Manners, J. C., Morcrette, C. J., Roberts, M. J., Stratton, R. A., Webster, S., Wilkinson, J. M., Willett, M. R., Boutle, I. A., Earnshaw, P. D., Hill, P. G., MacLachlan, C., Martin, G. M., Moufouma-Okia, W., Palmer, M. D., Petch, J. C., Scaife, A. A., and Williams, K. D.: The Met Office Unified Model Global Atmosphere 2.0/2.1 and JULES Global Land 3.0/3.1 configurations, Geoscientific Model Development, 4, 919–941, doi:10.5194/gmd-4-919-2011, 2011.

Table 1. Initial experiments comparing CICE under 6 different resolutions

Experiment	Vertical resolution	Timestep length	Coupling period length
1 (Hi-res 1S)	1 cm (100 layers)	1 second	1 second
2 (Low-res 1S)	25 cm (4 layers)	1 second	1 second
3 (Hi-res 1H)	1 cm (100 layers)	1 hour	1 hour
4 (Low-res 1H)	25 cm (4 layers)	1 hour	1 hour
5 (Hi-res 3H)	1 cm (100 layers)	1 hour	3 hours
6 (Low-res 3H)	25 cm (4 layers)	1 hour	3 hours

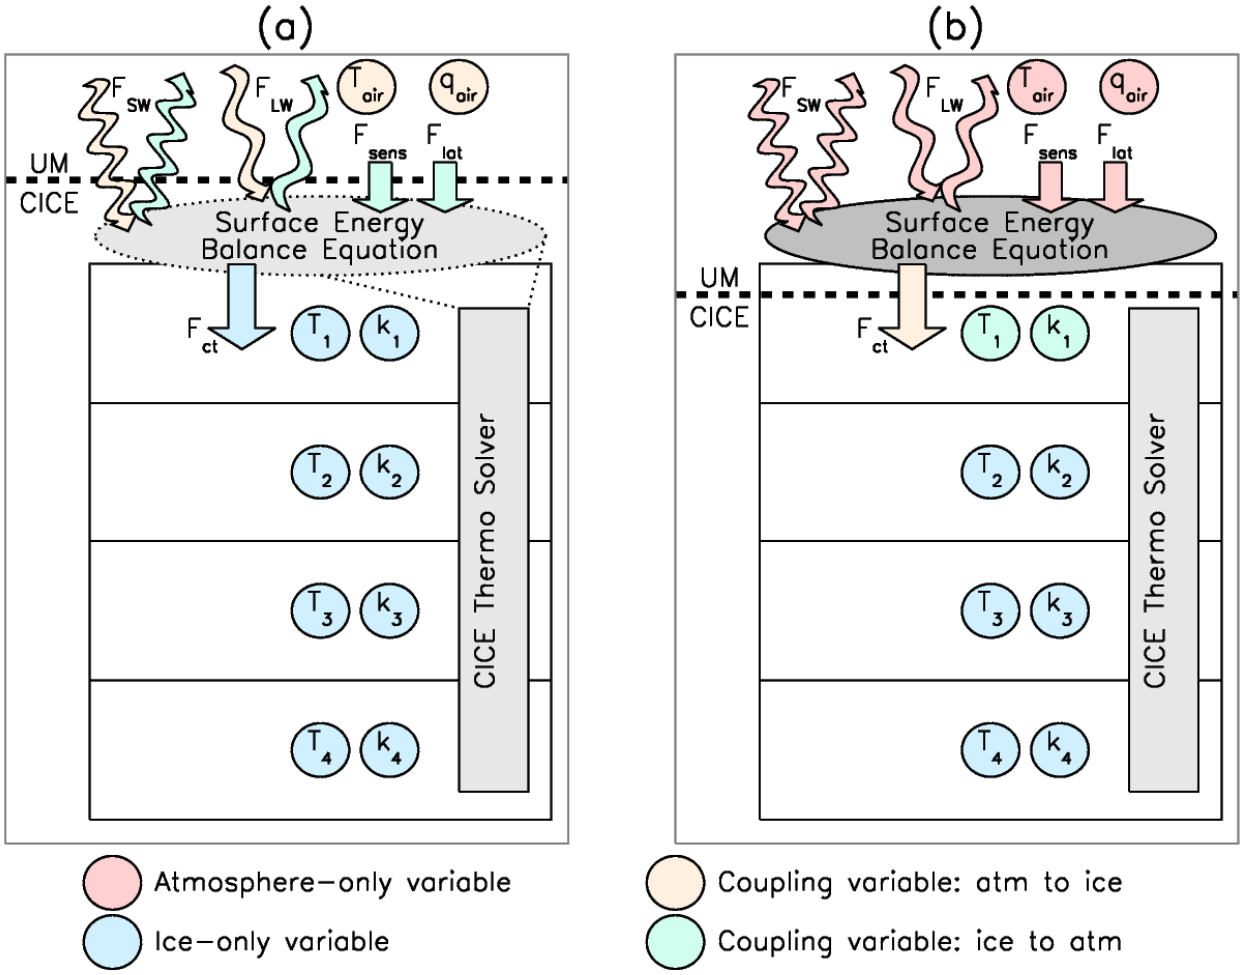

Figure 1: Schematics demonstrating (a) The 'CICE' coupling method; (b) the 'JULES' coupling method. In this figure F_{SW}, F_{LW}, F_{sens} and F_{lat} denote fluxes of shortwave radiation, longwave radiation, sensible and latent heat respectively, F_{ct} the conductive flux from the surface to the ice, T_{air} and q_{air} the temperature and humidity of the lowest atmospheric layer, and T_j and k_j, the temperature and effective conductivity of ice layer j.

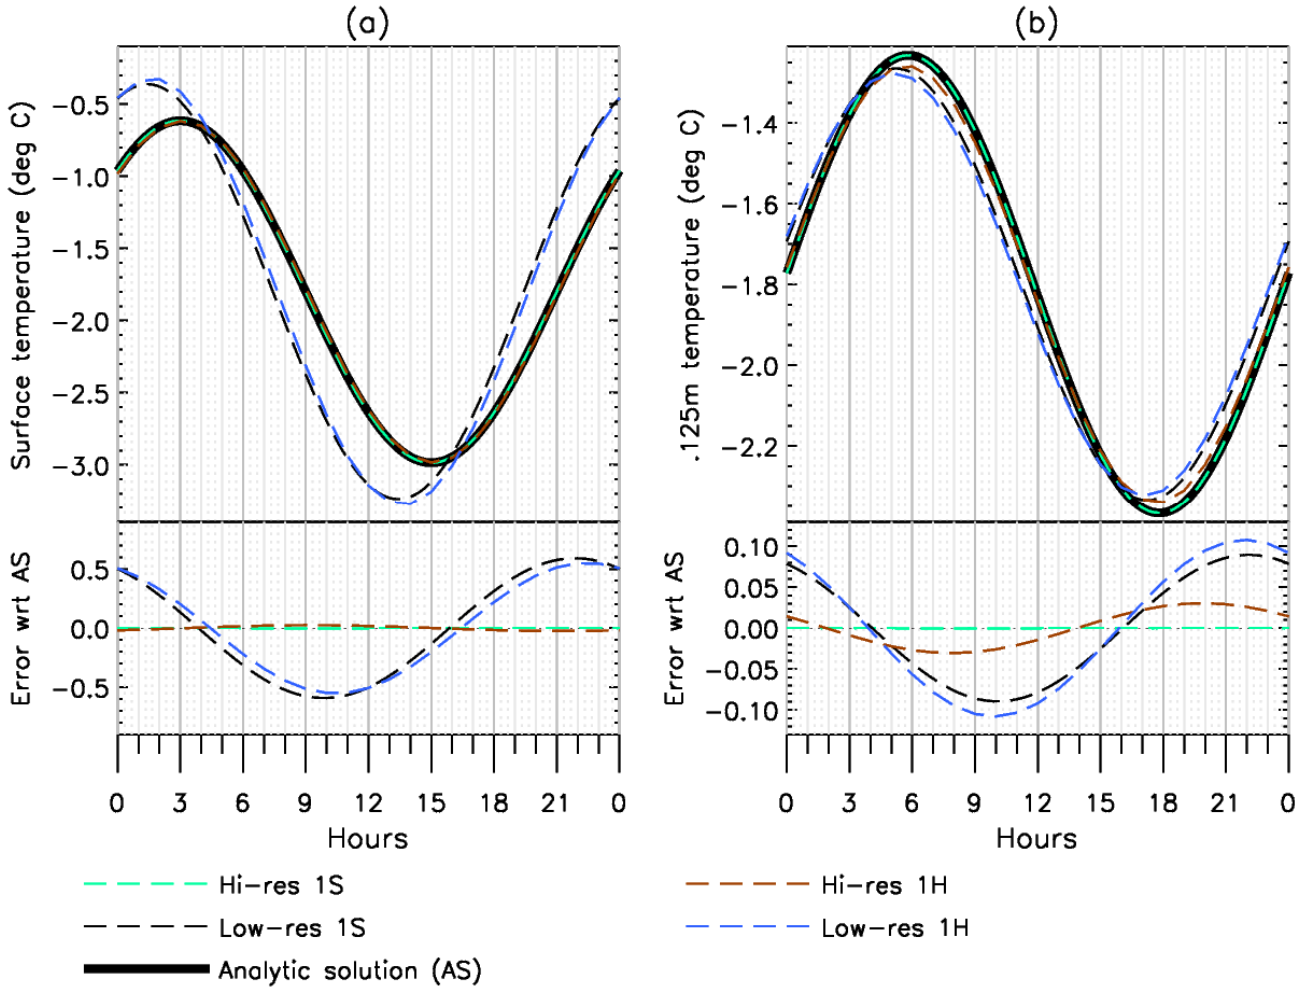

Figure 2: The performance of the CICE temperature solver under varying spatial resolution and timestep length, with coupling period 1 hour. Showing a) surface temperature; b) temperature at .125m depth.

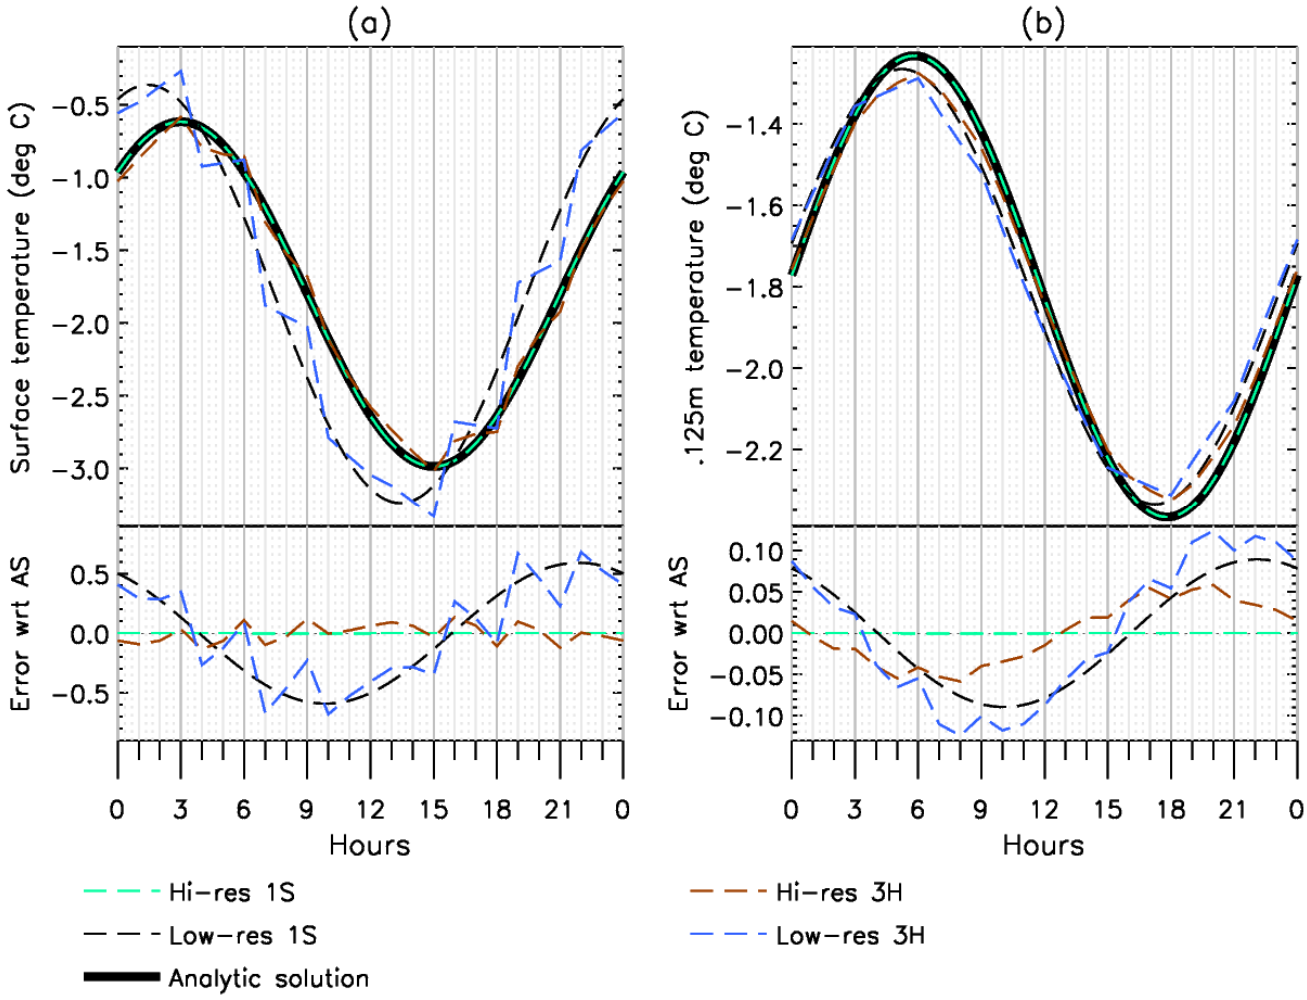

Figure 3: The performance of the CICE temperature solver under varying spatial resolution and timestep length, with coupling period 3 hours. Showing a) surface temperature; b) temperature at .125m depth.

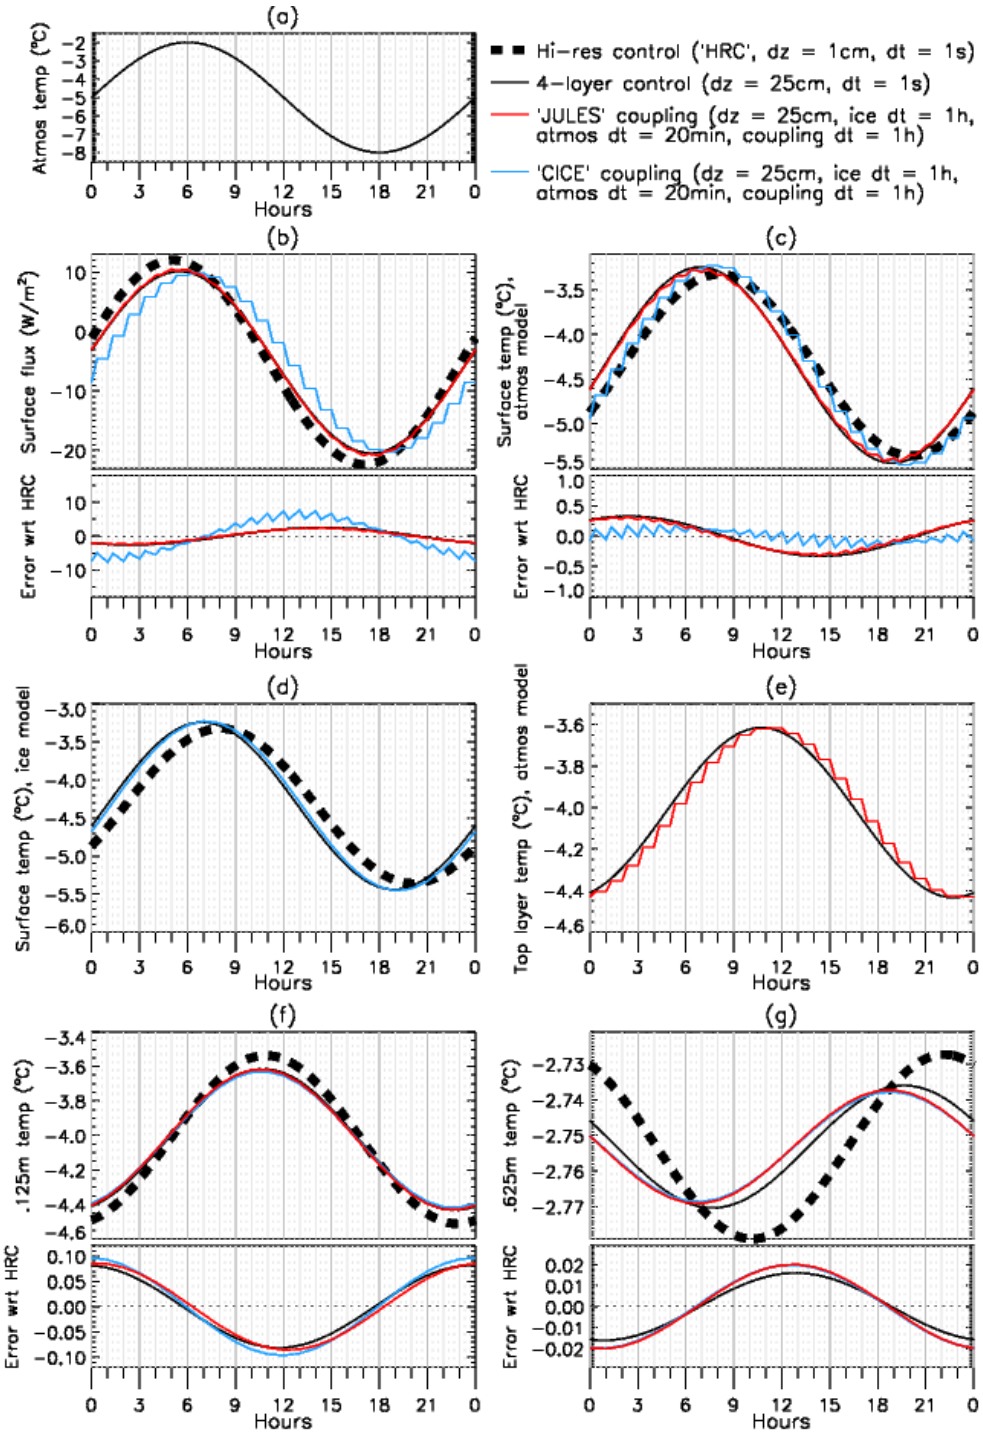

Figure 4: Comparing the two coupling methods, with a 1 hour coupling period. Showing a) atmospheric air temperature (the experiment forcing); b) surface flux; c) surface temperature, as seen by the atmosphere; d) surface temperature as seen by the ice; e) .125m ice temperature as seen by the atmosphere; f) .125m temperature as seen by the ice; g) .625m temperature.

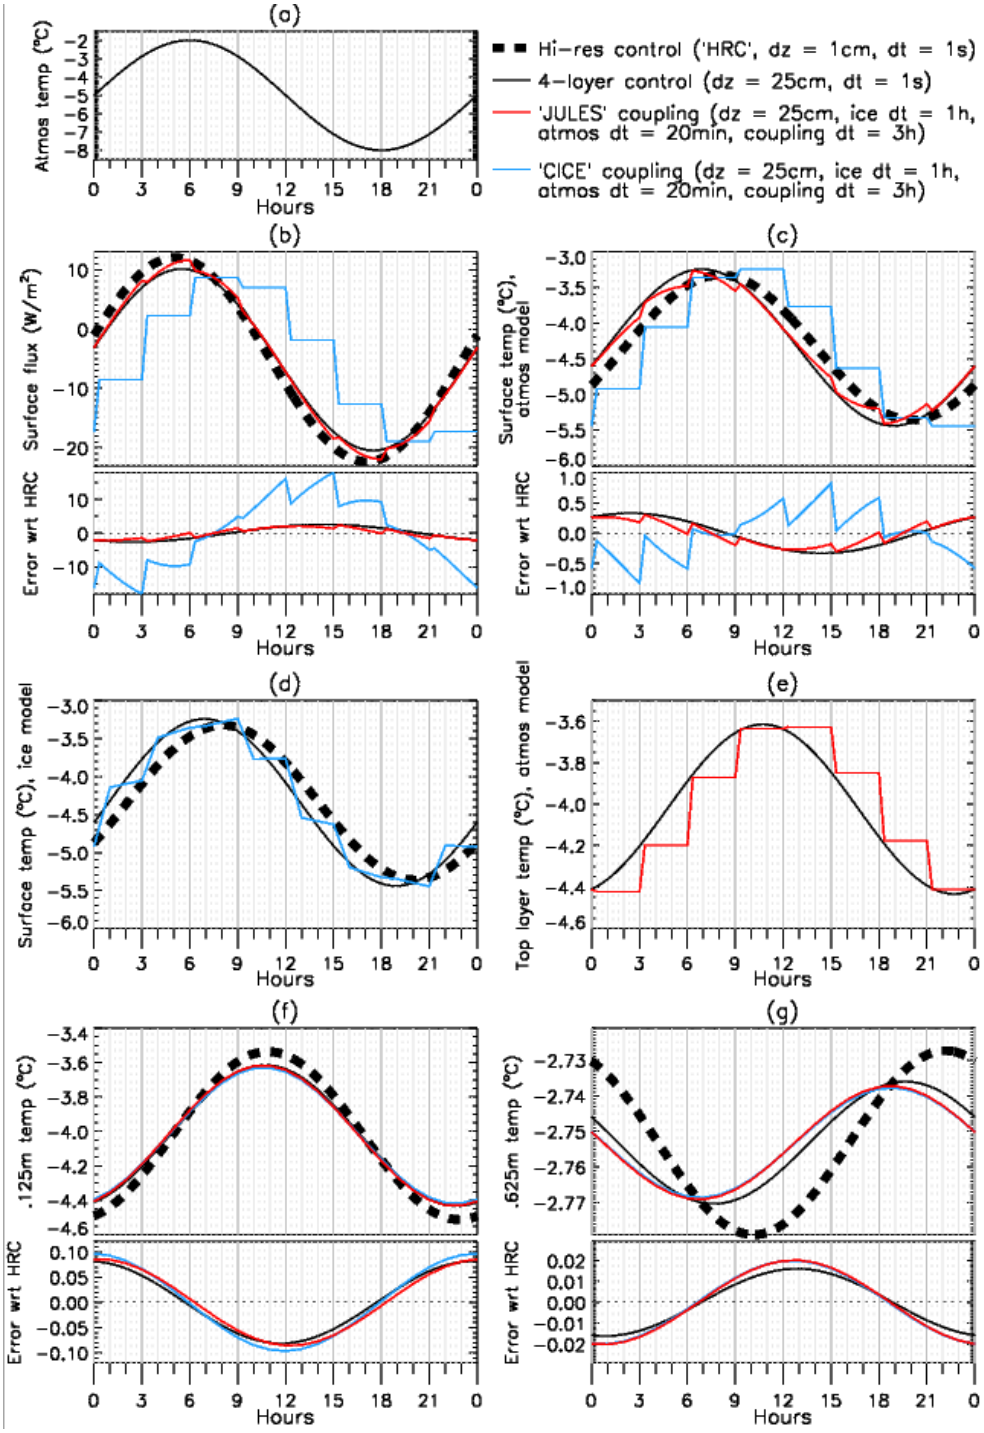

Figure 5: Comparing the two coupling methods, with a 3 hour coupling period. Showing a) atmospheric air temperature (the experiment forcing); b) surface flux; c) surface temperature, as seen by the atmosphere; d) surface temperature as seen by the ice; e) .125m ice temperature as seen by the atmosphere; f) .125m temperature as seen by the ice; g) .625m temperature.

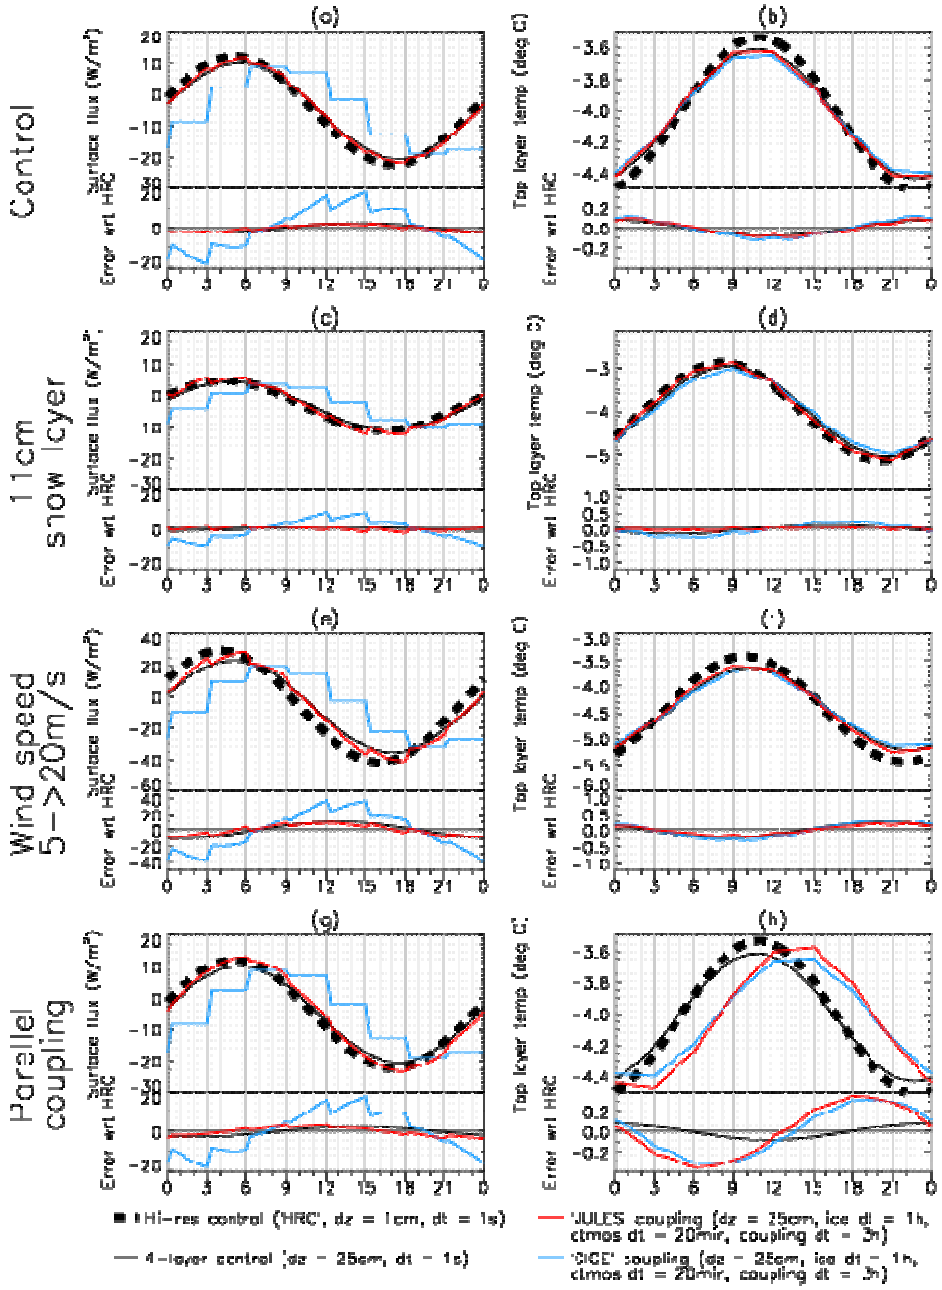

Figure 6: Demonstrating the performance of the two coupling methods when the parameters of the experiment are varied. Showing surface flux (left) and top-layer temperature (right) for (a,b) Original experiment; (c,d) an experiment in which an 11cm snow layer was present on the ice; (e,f) an experiment in which wind speed was increased from 5 m/s to 20 m/s; (g,h) an experiment in which parallel, rather than serial, coupling was employed. For the snow layer experiment, the top-layer temperature represents the temperature of the snow layer; for all others, it represents the temperature of the top ice layer.

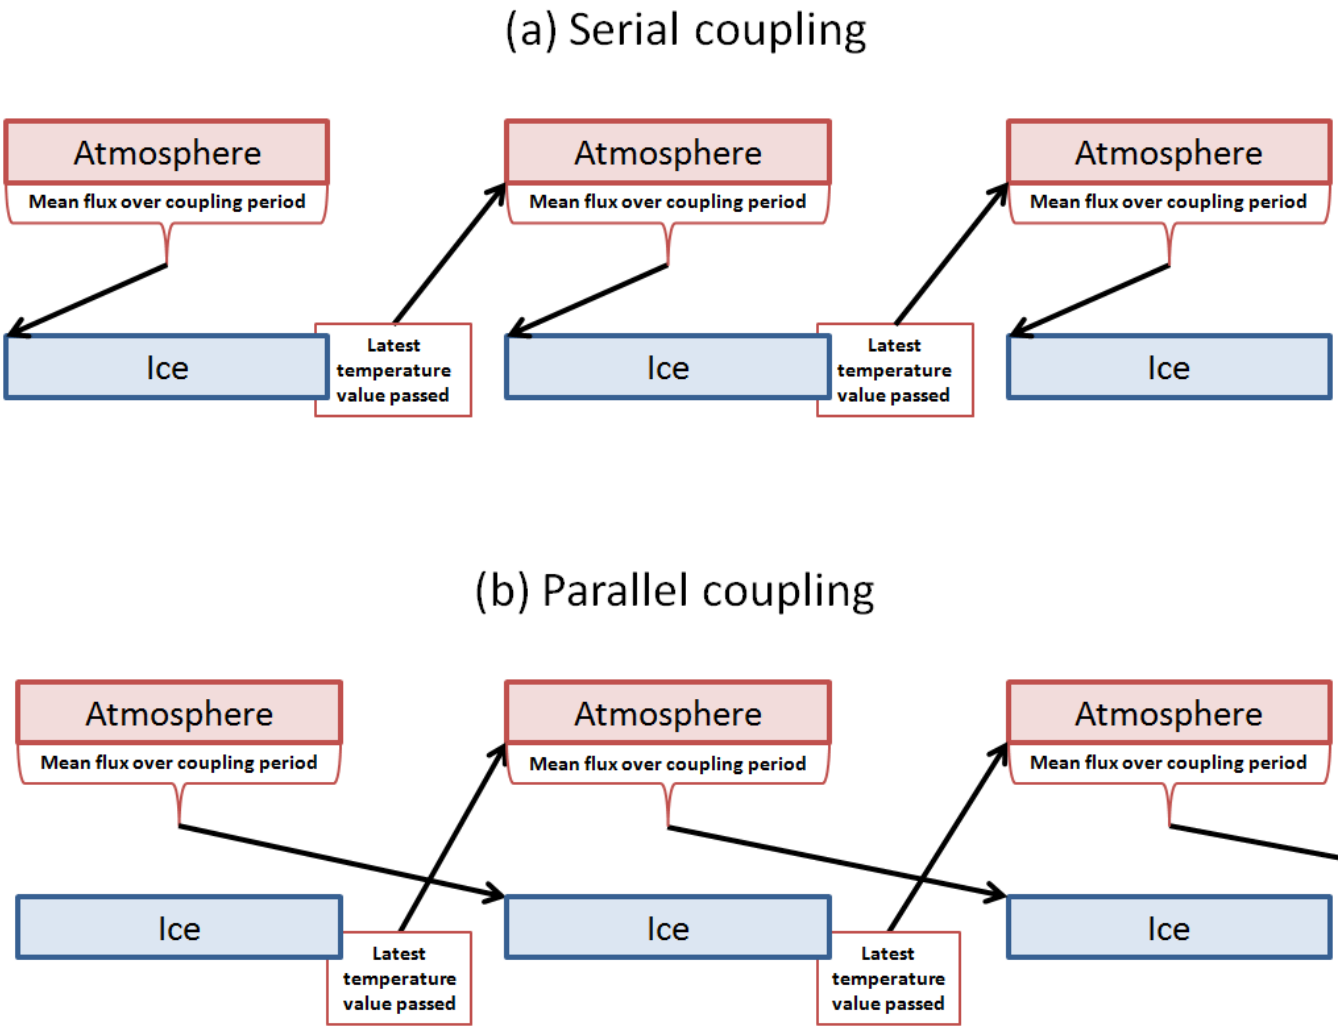

Figure 7: Schematic demonstrating the (a) serial and (b) parallel coupling frameworks, as described in Section 4.4.

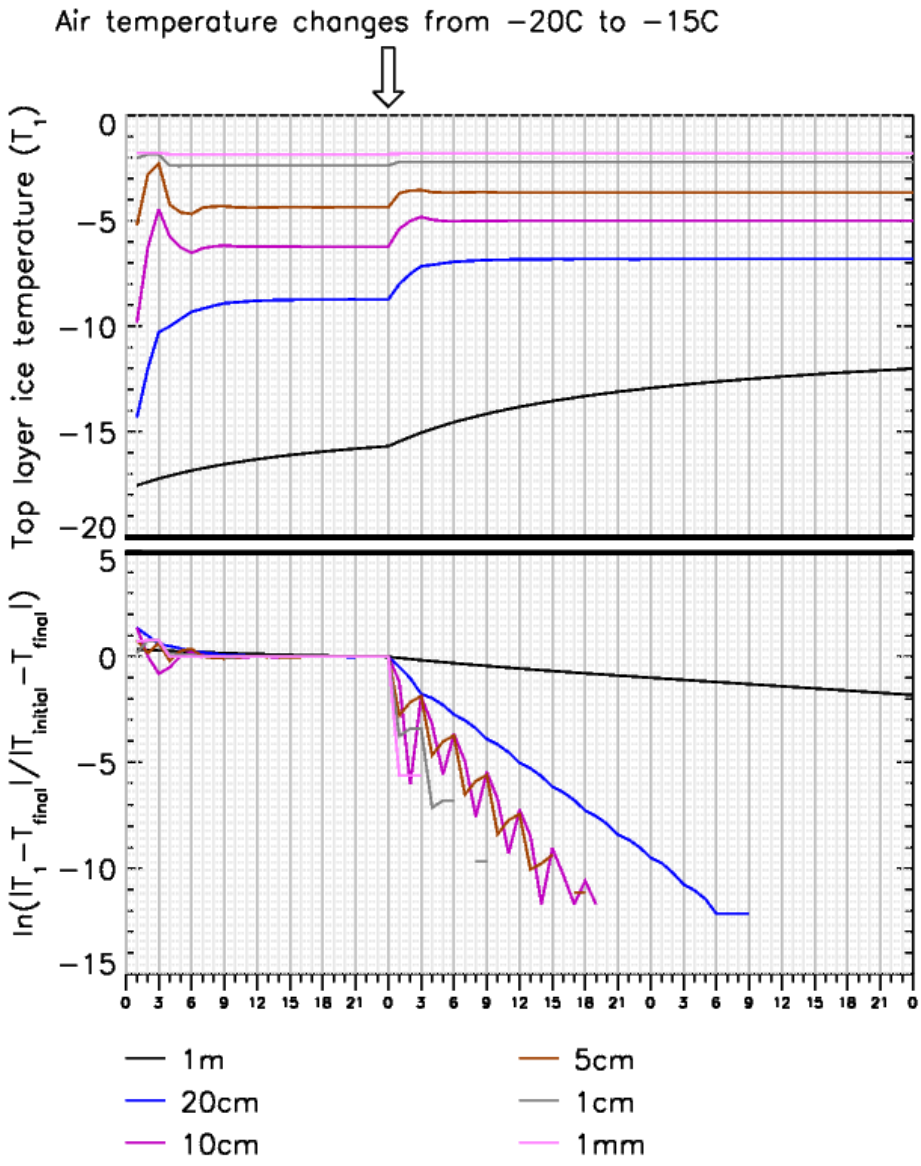

Figure A1: Showing the evolution of top layer ice temperature following a sudden change in air temperature, under the 'JULES' coupling method. The lower panel shows the evolution of $\ln\left(\frac{|T_1 - T_{final}|}{|T_{initial} - T_{final}|}\right)$ **to allow easy comparison of the rates of convergence for differing ice thicknesses, where** T_1, T_{final} **and** $T_{initial}$ **respectively refer to the evolving top layer ice temperature, the value of top layer ice temperature after 3 days, and the value at 1 day, at the time of the perturbation.. The graph 'disappears' when the difference falls below minimum precision.**

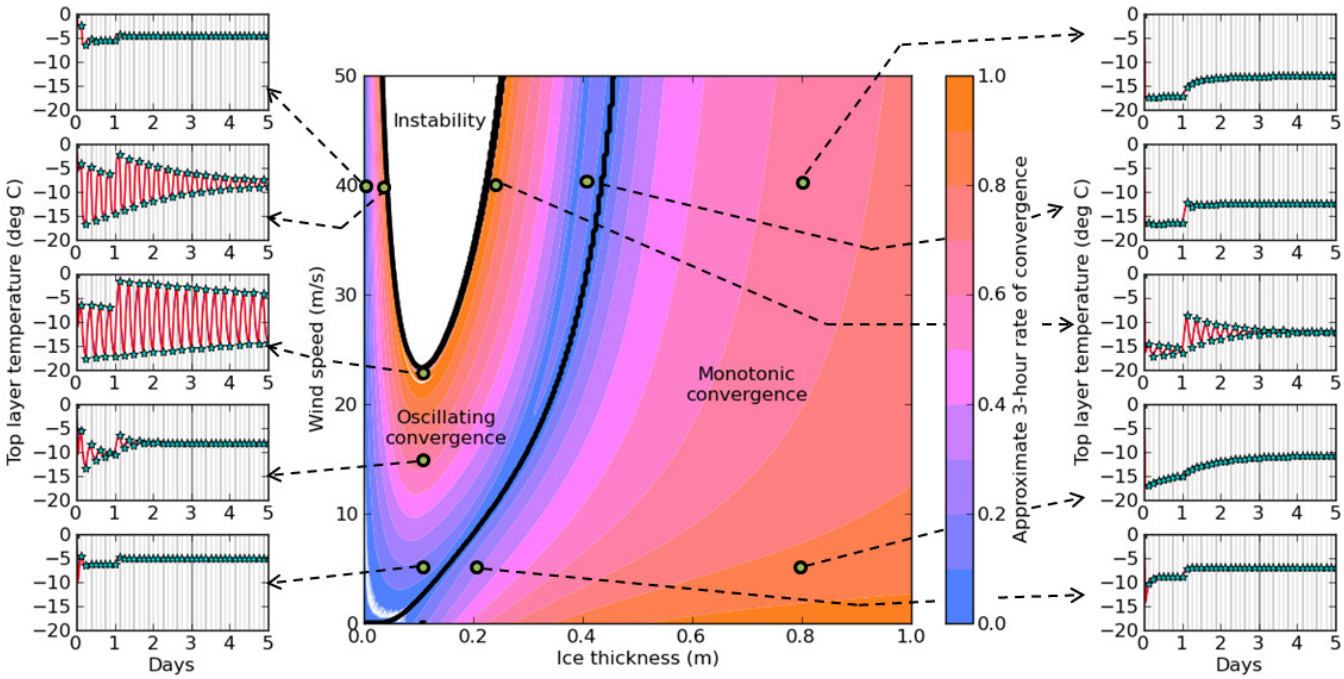

Figure A2: 'Map' of stability of the coupled ice and surface solvers, as ice thickness and wind speed are varied, with a 3-hour coupling period. Speed of convergence is indicated in colour, with blue = rapid convergence, red = slow convergence. Regions of 3-hour monotonic convergence, 3-hour oscillating convergence and instability are indicated. Timeseries of top layer ice temperature are shown for 10 representative points of the variable space.

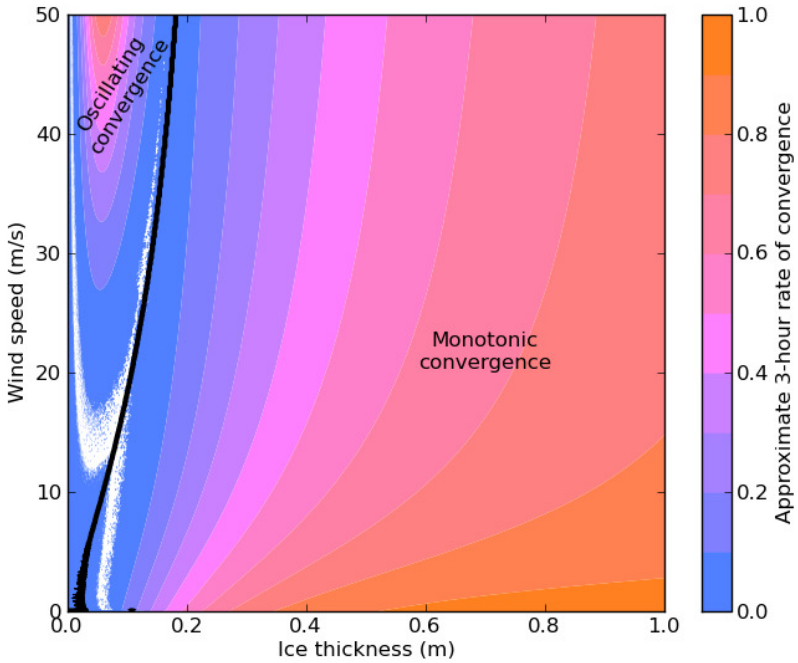

Figure A3: 'Map' of stability of the coupled ice and surface solvers, as ice thickness and wind speed are varied, with a 1-hour coupling period.