Comparison of the preindustrial climate of the IPSL-CM5A-LR model on different computers used at IPSL

J. Servonnat (1), M.A. Foujols (2), F. Hourdin (3), A. Caubel (1), P. Terray (4), O. Marti (1)

(1) Laboratoire des Sciences du Climat et de l'Environnement
(2) Institut Pierre Simon Laplace
(3) Laboratoire de Météorologie Dynamique
(4) Laboratoire d'Océanographie et de Climatologie : Approches Expérimentales et Numériques

The 15th of December 2012, the NEC SX-9 computer called "Mercure" was definitively shut down. This computer was dedicated to the IPSL climatic simulations since April 2009. It has been used to produce a great part of the IPSL simulations following the CMIP5 (*Coupled Model Intercomparison Project 5th phase*) protocol, notably for the science that was aimed at the IPCC AR5 (Intergovernmental Panel on Climate Change, Assessment Report number five). The promotion and use of those simulations in today's and future research is a major issue.

With Mercure being turned off, the climate modelers run the IPSL climate model on the other available computers: Titane at CCRT, Vargas and Ada at IDRIS, and Curie at TGCC. But can we take the model from one computer to another during an on-going study, with no consequence on the robustness of the results? In other words, is the climate of the model the same when ran on different computers?

To answer this question, it was necessary to compare the climate simulated by the IPSL model on different computers. As a reminder, the IPSL climate model is a computing program composed of numerous sub-programs representing hundreds of thousands of code lines. For each point of the grid that discretizes the terrestrial globe (in latitude, longitude, altitude or depth), the programs set the calculations to be done, for each time step, to quantify the different state variables of each component of the climate system (atmosphere, land surfaces, ocean, sea ice, biogeochemistry, atmospheric chemistry) after the equations describing their evolution. The current computers are composed of a great number of computing cores working in parallel which allow simulating the evolution of climate on period length ranging from several months to tens, hundreds, even thousands of years. Each processer (computing core) handles a subset of the total calculations, generally a geographical subdomain of the globe.

Our protocol is to compare simulations achieved with the same IPSL-CM5A-LR model version compiled on different computers, with the same boundary conditions, and starting from the same initial state. Five simulations were run this way (Figure 1): piControl2M0 on Mercure SX-9 (CCRT-TGCC), piControl2T0 on Titane (CCRT-TGCC), piControl2C0 on Curie (CCRT-TGCC), piControl2V0 on Vargas (IDRIS) and piControl2A0 on Ada (IDRIS). Two simulations are 100-year long (piControl2M0 and piControl2V0), and the others (piControl2T0, piControl2C0 and piControl2A0) are 200-year long (initially 100-year long but extended for the analyses, see further). We used the preindustrial climate control simulation named piControl2 as initial state and reference. This long and equilibrated 1900-year long simulation (1800-3699 period, dummy calendar) has been and still is used in several studies involving the IPSL-CM5A model. We took December of year 2499 as a starting point for our five simulations. In addition to those five simulations, we take an historical simulation (covering the true 1900-1999 period with the appropriate boundary conditions, such as the evolution of the greenhouse gases), called Historical in this study, to show a simulation that is actually very different from piControl2.

Figure 1 shows the global average of the temperature at 2m in the different simulations. On this figure we see that the different simulations do not follow the same trajectory. The reason for this behavior is the rounding differences on floating operations between the different computers. Those differences can even occur on the same computer because of a different compiler, compiling options or the splitting of the geographical domain for parallel computing. Those (though very small) differences act as analogues of infinitesimal perturbations at each time step. Because of the chaotic nature of the equations of meteorology, the differences between two simulations increase with time. However, different trajectories do not directly imply a different climate and we have to evidence whether the mean state of the climate of the model can be significantly different on the different computers compared with the SX-9. The piControl2M0 simulation has been ran to illustrate the influence of such small perturbations on the trajectory of the model: the code, computer, compiler, initial state and boundary conditions are the same as for piControl2, the only change is the file for the loading balance (Bands) for the different processers on Mercure.

To estimate whether two simulations are statistically different or not, we propose the following methodology: for a period starting at year 2500, of length T ranging from 20 to 200 years (by five-year increment), we calculate a quadratic distance *QDIST* (equation 1) between the climatology (2D spatial field) of piControl2 $X_{T,piControl2}$ calculated on this period and the climatology $X_{T,calc}$ calculated on the same period *T* for a simulation ran on the computer *calc*.

$$QDIST(T, calc) = \sum_i^N \left(w_i \times \left[X_{T,calc}(i) - X_{T,piControl2}(i) \right]^2 \right) \qquad \text{(eq.1)}$$

The index i corresponds to the grid points (N in total); w_i is a weight to take into account the area of the grid cells (associated with the grid points). The sum of the w_i is equal to one.

With the increase of the period length, we expect the climatology to converge toward the mean state of the climate model; if this mean state is the same on Mercure SX-9 (piControl2 reference simulation) and on another computer, *QDIST* should tend to zero with an increasing *T*. The results for the different computers are shown on Figure 2.

To have an estimate of the range of *QDIST* calculated between two climatologies considered as belonging to the same climate, we estimate the distribution of *QDIST* between two periods randomly sampled in piControl2, for each period length *T* (grey shaded area on Figure 2, see the caption for details on the method).

We present the results for three classic atmospheric variables used in climatology: temperature at 2m (t2m), precipitations (precip) and Sea Level Pressure (SLP). We also looked (not shown but commented in the text) at the radiative budget at the surface, at the top of the atmosphere and the surface pressure. The analyses are conducted on annual mean time series.

For t2m, we see on Figure 2a that all the simulations stay within the distribution of *QDIST* in piControl2 but Historical (which is straightforward because of the global warming of the 20th century). We even note that the simulations lie in the lower part of the distribution (below the 25th quantile) for period length greater than fifty years. This indicates that the climatologies of t2m of the different simulations achieved on the other computer than SX-9 are very close to the climatology of the model on SX-9. Similar results were obtained for the radiative budget at the surface and the top of the atmosphere.

For the precipitations (Figure 2b), the simulations ran on Titane, Vargas, Ada and the second simulation on SX-9 stay within the distribution of *QDIST* in piControl2. The simulation on Curie exceeds the 99th quantile of the *QDIST* distribution for some period lengths. This is also the case for the Historical simulation. We note however that for *T* greater than 70 years, the simulation on Curie and even Historical are within the distribution of the reference. The climatology of the precipitation of the model on the different computers tends to converge toward the climatology of the model on SX-9 when we increase the period length on which we calculate the climatology (i.e. when we use a more robust estimate of the mean state of the model).

For SLP, the Historical simulation comes out of the *QDIST* reference distribution for a period length greater than 95 years. The simulation on Titane is close to or greater than the 95th quantile for period length comprised between 70 and 115 years. Initially, the five simulations were only 100-year long; but this result of piControlT0 exceeding the 95th quantile for period length between 70 and 100 years, we decided to extend three of the simulations to see whether they could diverge from the piControl2 reference distribution for longer periods. The results we obtained actually confirmed that all the simulations converge within the *QDIST* reference distribution with the increase of the period length.

To investigate further the behavior of piControl2T0 for period length ranging between 70 and 115 years, we looked at the map of the difference between the SLP climatology (averaged on a 100-year period) of piControl2T0 et piControl2. This map shows a pattern with a negative pole over the Antarctic that closely reminds the Southern Annular Mode (SAM). The SLP time series within piControl2T0 averaged over this region shows substantial low frequency variability, characterized by

a continuous decrease until year 2550, and then a continuous increase to the end of the period (coming back to the level of 2500). This behavior of *QDIST* in piControl2T0 thus reflects the expression of the low frequency in the coupled model.
The results obtained for t2m and the radiative budget are comforting. Those for the precipitations on Curie may deserve further investigations, but the fact that *QDIST* converges within the reference distribution for period length greater than 115 years show that the precipitation climatology on Curie is not different from the one on SX-9. For SLP, all simulations converge toward the climatology of the model on SX-9 for long periods. However, it has been necessary to extend the simulations to be sure that the problem we saw in piControlT0 was only an issue of low frequency, and not a question of model drift toward a different mean state between Titane and Mercure.
This study has allowed us to lay the foundations of the comparison of the climate of the IPSL-CM5A-LR model on different computers used at IPSL. It also allowed us to see that the climatology of the frequently used atmospheric variables was not dependent of the computer used to run the model. We have also seen that the low frequency variability of the model was potentially substantial, and that it was useful to compare the simulations on periods longer than hundred years even for atmospheric variables. We can expect the need for even longer simulations for robust results on oceanic variables. After those first comforting results on the mean state of the model, the next step for the comparison of the climate of the model on different computers would be to assess the climatic variability of the model.

Acknowledgments
The authors thank Claude Frankignoul (LOCEAN), Francis Codron (LMD), Pascale Braconnot (LSCE) and Yann Meurdesoif (LSCE) for their suggestions and discussions. The computing has been done at CCRT, TGCC and IDRIS within the GENCI project. We here really want to acknowledge all the persons who fund, install, make evolve, fix the problems and help us to use the computing resources that are indispensable for our research projects.

Figures

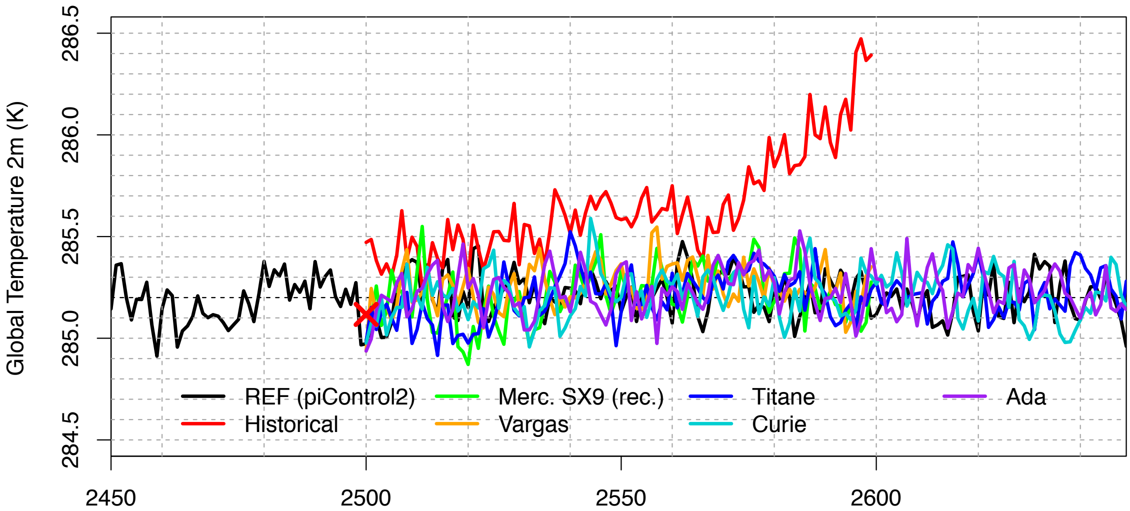

Figure 1: global mean temperature at 2m in piControl2 (solid black line, reference simulation on Mercure SX-9), piControl2M0 (in green, perturbed simulation on SX-9), piControl2V0 (in orange, on Vargas), piControl2T0 (in blue, on Titane), piControl2C0 (in light blue, on Curie) and piControl2A0 (in purple, on Ada). The historical simulation (Historical) is in red. The red cross highlights the common starting date for all simulations (but Historical).

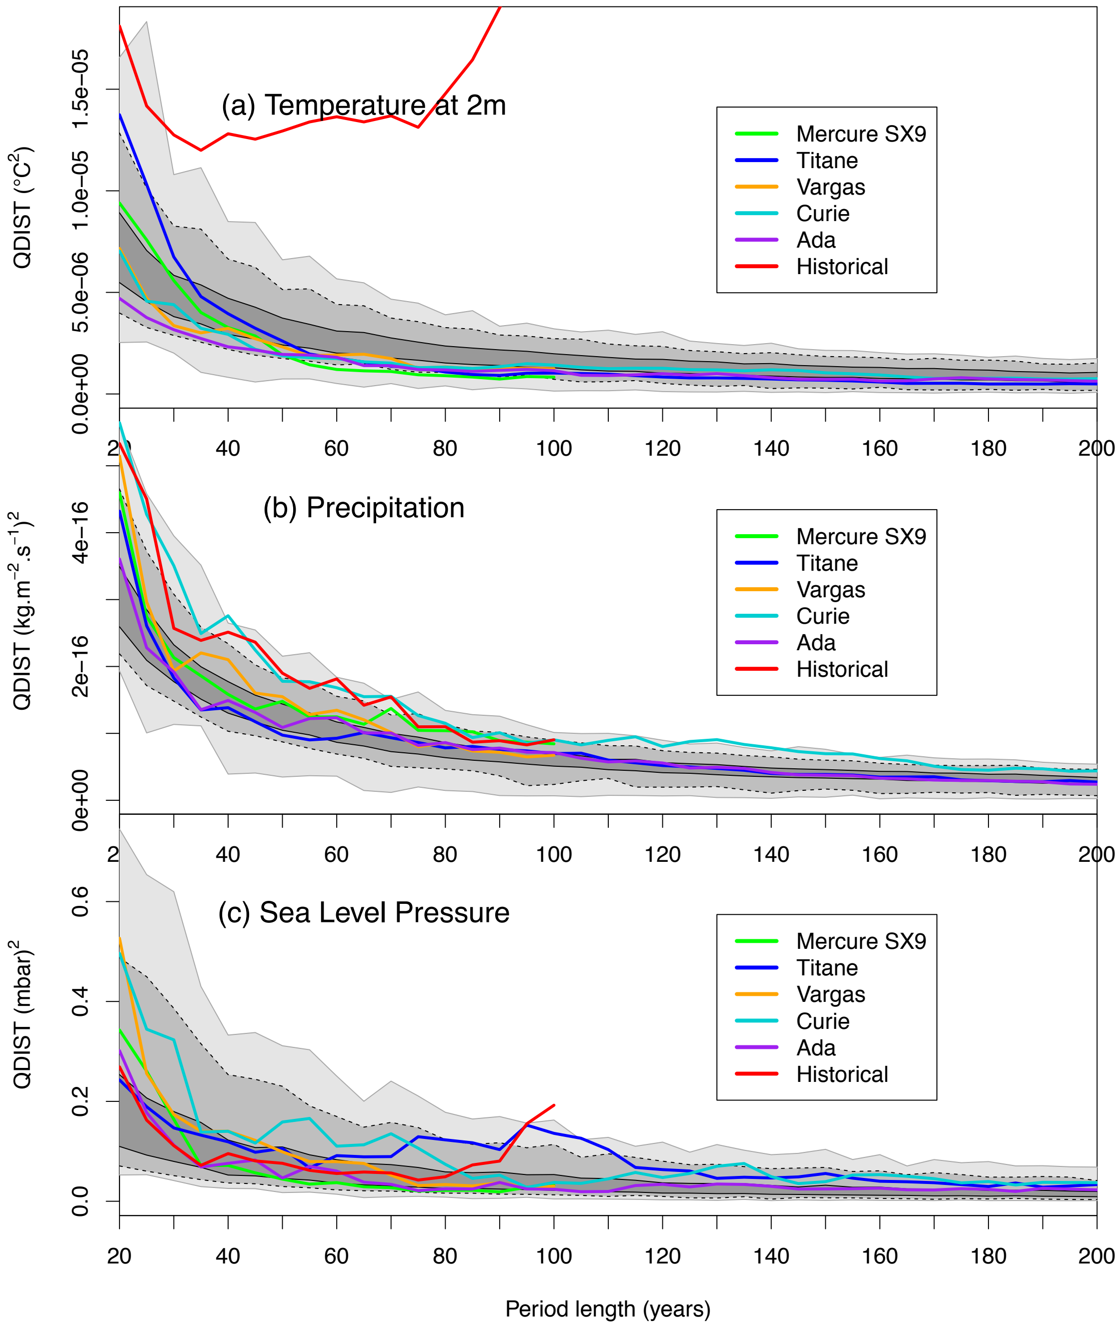

Figure 2: quadratic distance *QDIST* between the climatologies of the simulations achieved on the different computers (colored lines, color code in the legend box), the historical simulation, and the piControl2 reference simulation on the common 2500-2699 period. In abscissa is the period length (*T* in equation 1, in years) on which are calculated the climatologies. The upper panel (a) shows the results for the temperature at 2m, the middle panel (b) shows the results for the precipitations and the lower panel (c) the results for SLP. The grey shaded area shows the QDIST distribution calculated between two randomly taken periods in the 1900-year long piControl2 reference simulation. The starting dates of the randomly sampled periods (600 draws with replacement) are taken at a minimum of x years of interval for a climatology calculated on x years. From the lighter to the darker, the gray shading show the intervals 1/99th (highlighted by thin gray lines), 5/95th (highlighted by the dashed lines) and 25/75th (highlighted by the thin black lines) quantiles of the QDIST distribution between two randomly sampled periods in piControl2.