How to download, install, compile, and run OLAM-SOIL

Computer system Requirements

The size of computer platform required to run OLAM-SOIL depends on the scope and goals of the simulation. Some applications can be run at relatively low spatial resolution, implying that relatively few model grid cells are used, while others require high resolution with many cells. The time period represented in the simulation may be as short as hours or days, or as long as decades or centuries. The specified duration of the simulation divided by the length of a numerical time step used in the model (typically 1 minute), determines how many time steps the model will run. The computational work load is roughly proportional to the product **P** of the number of grid cells and the number of time steps. If **P** is below about 10^{10}, the simulation can normally be run on a single computer core (such as a laptop computer) in a day or less. However, many of today's cutting-edge simulations are much more complex and involve a **P** of 10^{15} or more. These simulations require a large parallel computing system of several thousand computer cores.

Anticipating that OLAM-SOIL modelers are likely to require a large computing system for at least some simulations, the developers of OLAM-SOIL have focused their resources to achieve high parallel efficiency on these systems. The model specifically targets large distributed-memory clusters of multi-core (cpu) shared-memory computing nodes with a Linux or Unix operating system, which is the most common environment used in large university and research laboratory computing platforms. The following software is required to run OLAM-SOIL with its full set of options and capabilities:

Linux/Unix operating system (e.g., Ubuntu, CentOs, Red Hat)
Fortran 90 compiler (Intel strongly recommended)
C compiler (gcc is good)
HDF5 (the data format used for most I/O)
MPI
NCAR Graphics (also called NCL)

Most university and research laboratory computing systems will already have installations of the above software libraries, but for any libraries that are absent, a system administrator can perform the installation. HDF5 and MPI should be compiled with Fortran 90 support. When HDF5 is compiled with Fortran support, it installs a script called h5sfc or h5pfc that calls the Fortran compiler and links in the proper HDF5 library and Fortran modules. This script is referenced in an OLAM-SOIL file, which will be described in the next section.

NCAR Graphics is less widely used than the other libraries in the above list and is least likely to be installed on the system. The easiest way to install NCAR Graphics is to download a pre-compiled binary (executable) file that was compiled under the same Linux operating system as

used on your computer. Binaries and instructions for this procedure can be found at https://www.ncl.ucar.edu/Download/ and https://www.ncl.ucar.edu/Download/install.shtml. Since the binary file is probably compiled with gfortran, it is necessary to install gfortran on your computer in addition to the Intel Fortran compiler used to compile the model.

NCAR Graphics is the primary graphics package that OLAM-SOIL uses to plot a wide variety of quantities generated by a model simulation. OLAM-SOIL contains a large library of basic prognosed quantities and secondary derived quantities that can be plotted, as well as a selection of add-ons (such as model grid cells, geographic maps, and wind vectors), that use the NCAR Graphics interface. However, some modelers have preferred to use other graphics packages such as MATLAB to plot and analyze OLAM-SOIL output, and it is possible to build and run OLAM-SOIL without NCAR Graphics at all. A standard interface to MATLAB is not yet available with OLAM-SOIL, although it is under development.

Many users will find it convenient to install OLAM-SOIL on small single-node or shared-memory computers for small test runs and low-resolution simulations. For these systems, MPI is not required. However, the remaining software (Linux operating system, Intel Fortran 90 compiler, gcc C compiler, HDF5, and NCAR Graphics) is still required or strongly recommended.

Some have attempted to use gfortran to compile OLAM-SOIL, and to our knowledge this has not been successful due to documented bugs in the gfortran compiler. The last time we checked on this, the bugs had not yet been repaired. Years ago, we used Portland Group (PGF) Fortran 90 to compile a more basic version of OLAM, but the model did not execute as fast as when compiled with Intel Fortran 90. We have not used PGF in recent years and do not know if it would compile the model successfully. This is why we strongly recommend Intel Fortran 90 as the compiler of choice.

There have been recent attempts to build OLAM-SOIL on Windows computers (e.g. laptops) using the Cygwin version of Linux. At this time, we are unaware whether these attempts have been successful. Thus a word of caution that we cannot provide a roadmap for this procedure.

Note to those attending the April 8, 2018 OLAM-SOIL workshop at EGU:

During part of the workshop, we will be demonstrating how to run OLAM-SOIL with a small practice simulation. For those who want to perform the exercise on their own computers, it will be necessary to have OLAM-SOIL installed and compiled before the workshop, because there is insufficient time for this during the workshop. If you do not have a laptop with the recommended software, we encourage you to work with your system administrator to install and build the model on your university/laboratory computer so that you can perform a run remotely from your laptop during the workshop.

Downloading, installing, and compiling OLAM-SOIL

1. Create a directory called olam on your computer. Download the tarfile called olam_5.3.tar from the OLAM-SOIL website http://olam-soil.org/modeling/ and copy to your olam directory. Untar the file by giving the command 'tar xvf olam_5.3.tar'. This will create a subdirectory 5.3, which is the model version designation, and the following directories beneath it:

build_SOIL6
build_test
CMAQ
convect
ED2
etc
hurricane
leaf
MEGAN
modules
oisan
omodel
outils
radiate
sea
SOIL
test_cases

Most of these directories contain the model source code files, but the directories build_test and build_SOIL6 are template directories for compiling and running the model. The directory build_test is used for building the standard version of OLAM (with the LEAF4 land surface model) and the directory build_SOIL6 is used for building the new OLAM-SOIL version of the model. In normal circumstances, you would want to copy one or the other of these directories and its contents to a new directory (e.g., cp –R build_SOIL6 build_SOIL7) so that you can work in the new directory and leave the original template directory unchanged. Following the example of creating the build_SOIL7 directory, we will work in that directory.

2. Use an editor to open the file called include.mk in the build_SOIL7 directory. Ignore the upper third of this file, which contains commented-out lines (i.e., those beginning with the # symbol) that pertain to the PGF compiler (which you hopefully are not trying to use). The middle third of the file is for the Intel compiler, which is what we recommend that you use. Edit lines in this section so that they indicate to the model where to find the required software libraries on your computer. To indicate the compiler, set F_COMP=ifort or F_COMP=h5pfc. The ifort example points only to the compiler itself, while the h5pfc example points to the wrapper that executes the fortan compiler and also loads the hdf5 library; h5pfc is what you actually want to use if hdf5 has been compiled with fortran support as recommended. Follow comments in the include.mk file to guide you in setting IEEE_ARITHMETIC, OLAM_MPI, OLAM_PARALLEL_HDF5, PAR_INCS, PAR_LIBS, F_OPTS, FIXED_SRC_FLAGS, SLOW_FFLAGS, C_COMP, C_OPTS, NCARG_DIR, LIBNCARG, LOADER, and LOADER_OPTS. Many of these will

not need to be changed from the settings in the original file, but others will depend on the specific installation of the software libraries on your computer. The variables HDF5_LIBS and HDF5_INCS should be left empty when using the HDF5 wrapper script "h5fc" or "h5pfc". HDF5_LIBS and HDF5_INCS are only needed if F_COMP has been set to a Fortran compiler (such as "ifort"), in which case you will have to manually set the necessary HDF5 libraries and the location of the HDF5 Fortran modules. Set the NCARG_DIR path to where NCAR Graphics is installed on your computer. Note that the library links specified under LIBNCARG must include the gfortran library since gfortran was used to compile the NCAR Graphics binary (as recommended in the preceding section). This completes the required edits to include.mk.

3. Compile the model by simplying typing "make". This executes the Makefile in your build_SOIL7 directory. The Makefile makes use of other files in the directory, namely include.mk which you edited as per instruction 2, paths.mk, and sources.mk. Sources.mk, together with paths.mk, point to each of the source code files in the various directories that will be compiled. Compilation produces the executable file olam-5.3-mpi, or simply olam-5.3 if you removed 'yes' from the OLAM_MPI specification in instruction 2. Note: If NCAR Graphics is not installed on your computer (in which case you will presumably be plotting on a separate computer with NCAR Graphics or have interfaced OLAM-SOIL to a different graphics package), open the sources.mk file in an editor and change 'o_ncar.f90' to 'o_ncar_dummy.f90' so that no actual NCAR Graphics library code will be referenced in the model. Make this change before executing the 'make' command.

4. You can run the model (without MPI) simply by typing the executable file name on the command line: olam-5.3-mpi (or olam-5.3). If you are running on a larger computer with a batch job system, you can also execute the model in batch mode.

When the model begins execution, the first thing it does is to read the input namelist file, which has the name 'OLAMIN' and is located in the build_SOIL7 directory. The namelist variables in this file need to be modified by the user as appropriate to how the model is to be used for a given application. What each variable does and how it should be set is explained in detail in the OLAMIN User's Guide, a pdf file located in the etc subdirectory of OLAM.

Unless the user has set the OLAMIN variables for a very idealized simulation, OLAM will require that various geophysical datasets be available for input. These datasets include soil composition and other properties, vegetation class, sea surface temperature, sea ice abundance, topographic height, normalized difference vegetation index (NDVI), water table depth, and gross primary production (GPP) and total approximately 10 GB in size. These datasets can be downloaded from the OLAM-SOIL website http://olam-soil.org/modeling/. When installing OLAM, the user should create a directory called 'olamdatah5' directly beneath the top-level olam directory, and copy the geospatial data directories to the olamdatah5 directory. Default specifications in OLAMIN are set up to access these datasets from that location.