Running Distributed Schrödinger Jobs

A number of Schrödinger products can distribute work over multiple processors. There are several algorithms for performing the distribution. Some jobs divide the input structures into batches. Each batch is then submitted to a processor for execution as a subjob. The number of processors used and the number of batches (subjobs) can be different. Others divide the work among the available processors, so that the number of subjobs and the number of processors is the same.

Jaguar and QSite can use multiple processors, but use MPI parallel processing rather than distributed processing. Desmond uses MPI parallel processing, but also distributes multiple simulations.

For each distributed job there is a driver that is responsible for dividing up the work, submitting the subjobs, and collating the results. Thus, if *N* processors are requested, the job creates *N*+1 processes. Depending on the program, the driver can run on the local host or on the remote host. If the driver runs on the remote host, there are two separate cases to consider:

- The remote host (or collection of remote hosts) does not involve a queuing system. Even though the driver does not takes much time, the overall job does not necessarily run inefficiently, because the operating system can swap the driver process out when it is idle and use the time for other processes.
- The remote host is running a queuing system. The driver is run as a separate queued job, and the subjobs are then submitted to the queue by the driver. The driver job can then occupy a processor that will be idle most of the time. The exception is Prime loop refinement and Phase, for which the driver runs a subjob locally.

The hosts on which the driver and the subjobs run are set with one or more of the following options:

- `-HOST`: General list of hosts for the job. If one of other options is not used, the specified hosts are generally used for the other. If neither of the other options is used, the driver runs either on `localhost` or the first host, the subjobs can run on the remaining hosts or on all the hosts, including the first.
- `-DRIVERHOST`: host on which to run the driver (must be a single host name). If omitted, the driver may be run on the first host specified by `-HOST` or on `localhost`, if `-HOST` is not used or the application default is to run the driver locally.
- `-SUBHOST`: hosts on which to run the subjobs. If omitted, all hosts specified by `-HOST` are used for the subjobs. The number of CPUs used is determined by the number of hosts or host/CPU combinations specified.

The syntax of these options is given in Section 2.3 of the *Job Control Guide*. Information on the number of subjobs, number of CPUs, and the driver job location is given in Table 1 for running distributed jobs from the command line, and in Table 2 for running distributed jobs from Maestro. The notation "Standard" means that the options listed above are used to determine the driver location and the number of CPUs.

Table 1. Distributed processing behavior for jobs submitted from the command line.

Product	Program	Number of subjobs determined by	Number of CPUs for subjobs determined by	Driver location
Epik	`epik`	`-NJOBS` *njobs*, `-JOBCTS` *maxstructs*	Standard	Standard
Desmond	`multisim`	`-maxjob` *njobs*	`-cpu` *cpu-spec* or `.cfg` file; value is per subjob.	Standard
Glide	`glide`	`-NJOBS` *njobs*	Standard	Standard
Induced Fit Docking	`ifd`	Number of ligands (Glide), number of poses (Prime).	`-NGLIDECPU`, `-NPRIMECPU`, or keywords in input file	Standard
Ligand & Structure-Based Descriptors	`lsbd`	Keywords in input file	Keywords in input file	Local host
LigPrep	`ligprep`	`-NJOBS` *njobs*, `-JOBCTS` *maxstructs*	Standard	Standard

Table 1. Distributed processing behavior for jobs submitted from the command line. (Continued)

Product	Program	Number of subjobs determined by	Number of CPUs for subjobs determined by	Driver location
MacroModel	`bmin`	`-NJOBS` *njobs*, `-JOBCTS` *maxstructs*	Standard	Standard
MCPRO+	`mcpro_ddg`	Number of lambda points × 2	Number of hosts in `-HOST` list	Local host
	`mcpro_fep`	Number of lambda points	Number of hosts in `-HOST` list	Local host
	`mcpro_lrm`	Number of ligands	Number of hosts in `-HOST` list	Local host
Phase	`phase_dbsearch` `phasedb_confsites` `phasedb_convert`	Number of processors	Number of hosts in `-HOST` list	First host in `-HOST` list; `-DRIVERHOST` ignored.
Prime	`multirefine`	Stage of process, `MAX_JOBS` keyword in input file	Host given by `HOST` keyword in input file, otherwise host given by `-HOST`	First host in `-HOST` list
QM-Polarized Ligand Docking	`qpld`	`-NJOBS` *njobs*	Number of hosts in `-HOST` list or `-host_`*program* lists.	Host specified by `-DRIVERHOST`; default first host in `-HOST` list
SiteMap	`sitemap`	Number of processors	Number of hosts in `-HOST` list	Local host
Virtual Screening Workflow	`vsw`	`-NJOBS` *njobs*, `-adjust`	Number of hosts in `-HOST` list or `-host_`*program* lists.	Host specified by `-DRIVERHOST`; default first host in `-HOST` list

Table 2. Distributed processing for jobs submitted from Maestro.

Product	Job type	Number of subjobs set in or determined by	Number of CPUs set in or determined by	Master job location
CombiGlide	Docking	Start dialog box	Start dialog box	Local host
Epik		Number of CPUs (passed as `-NJOBS`)	Start dialog box	Selected host
Glide	Docking	Start dialog box	Start dialog box	Selected host
Induced Fit Docking		Number of ligands (Glide), number of poses (Prime).	Induced Fit Docking panel, Job options section.	Local host
Liaison	Simulation	Equal to number of CPUs	Systems folder, Liaison panel	Local host
LigPrep		Number of CPUs (passed as `-NJOBS`)	Start dialog box	Selected host
LSBD	Liaison	Number of CPUs	Advanced Options dialog box	Local host
	Prime MM-GBSA	Advanced Options dialog box	Advanced Options dialog box	
MCPRO+	Relative Binding Affinity by FEP	Number of lambda points × 2	Start dialog box	Local host

Table 2. Distributed processing for jobs submitted from Maestro.

Product	Job type	Number of subjobs set in or determined by	Number of CPUs set in or determined by	Master job location
	Free Energy Difference by FEP	Number of lambda points	Start dialog box	Local host
	Linear Response	Number of ligands	Start dialog box	Local host
Phase	Clean Structures, Generate Conformers, Find Matches	Number of CPUs	Start dialog box	Selected host
Prime	Loop Structure	Depends on stage of process and sampling method	Refine Structure - Options dialog box	Selected host
QM-Polarized Ligand Docking		Number of CPUs	Start dialog box	Selected host
SiteMap		Number of CPUs	Start dialog box	Selected host
Virtual Screening Workflow		Start dialog box	Start dialog box	Selected host

The launch directory must be available from the master job host for Induced Fit Docking, Liaison and LSBD.