

Defining wallclock time in the jobscript

Backfill scheduling is configured on all Pawsey's HPC systems. In a nutshell, backfill scheduling allows other jobs to use the reserved job slots, as long as it do not delay the start of jobs awaiting the reservation. In other words, backfill scheduling can be very useful for users running relatively shorter jobs. To make efficient use of this feature, you should specify the wallclock time of jobs as accurately as possible. If you know the approximate time that your job will run for, then setting it (with a suitable margin of error) in the `#SBATCH --time` line of your job script will enable the scheduler to make better choices for backfilling "gaps" in the machine.

When jobs finish, the resources assigned are then available to the next job. At Pawsey the next job is usually the highest in priority. If not enough resources are available for the next job, the resources will be held and accumulated until the job can start. Since all jobs have a maximum walltime, SLURM knows the worst case time for when the highest priority job can start. In the case of reservations (such as for maintenance), SLURM knows exactly when it will start. Holding up resources produces gaps where there is no job running, in which a small enough job may fit without causing a delay to the expected start time of the highest priority job.