MPI support
In Parsl and Globus Compute

Yadu Babuji
(Parsl and Globus Compute team)

MPI, why is it so hard!

ParslFest 2023

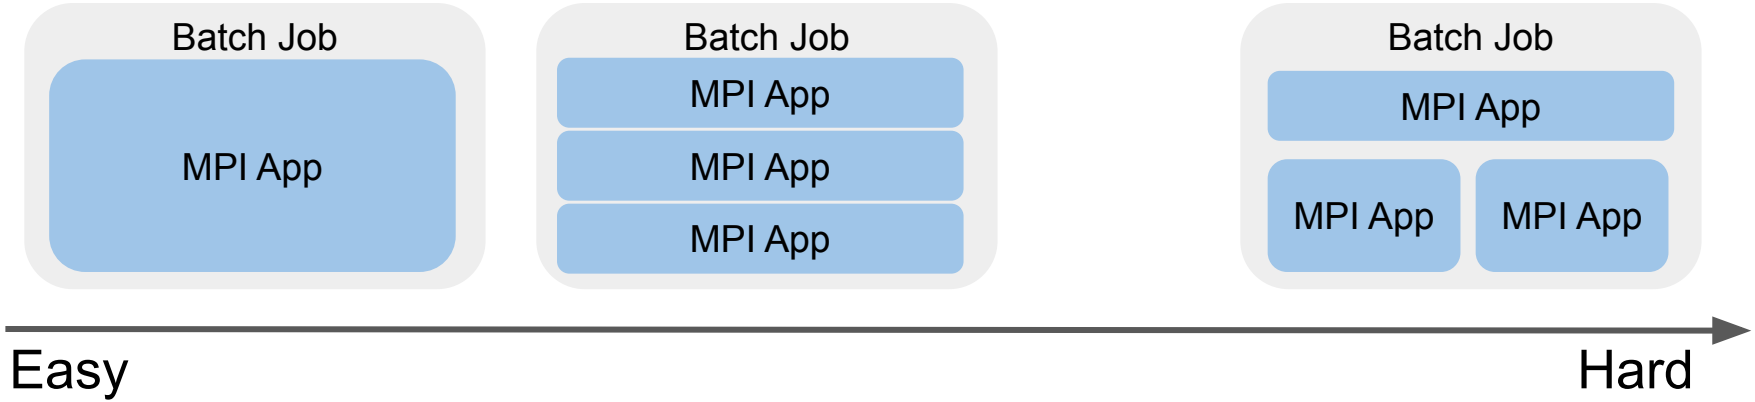

MPI support is here in Parsl!

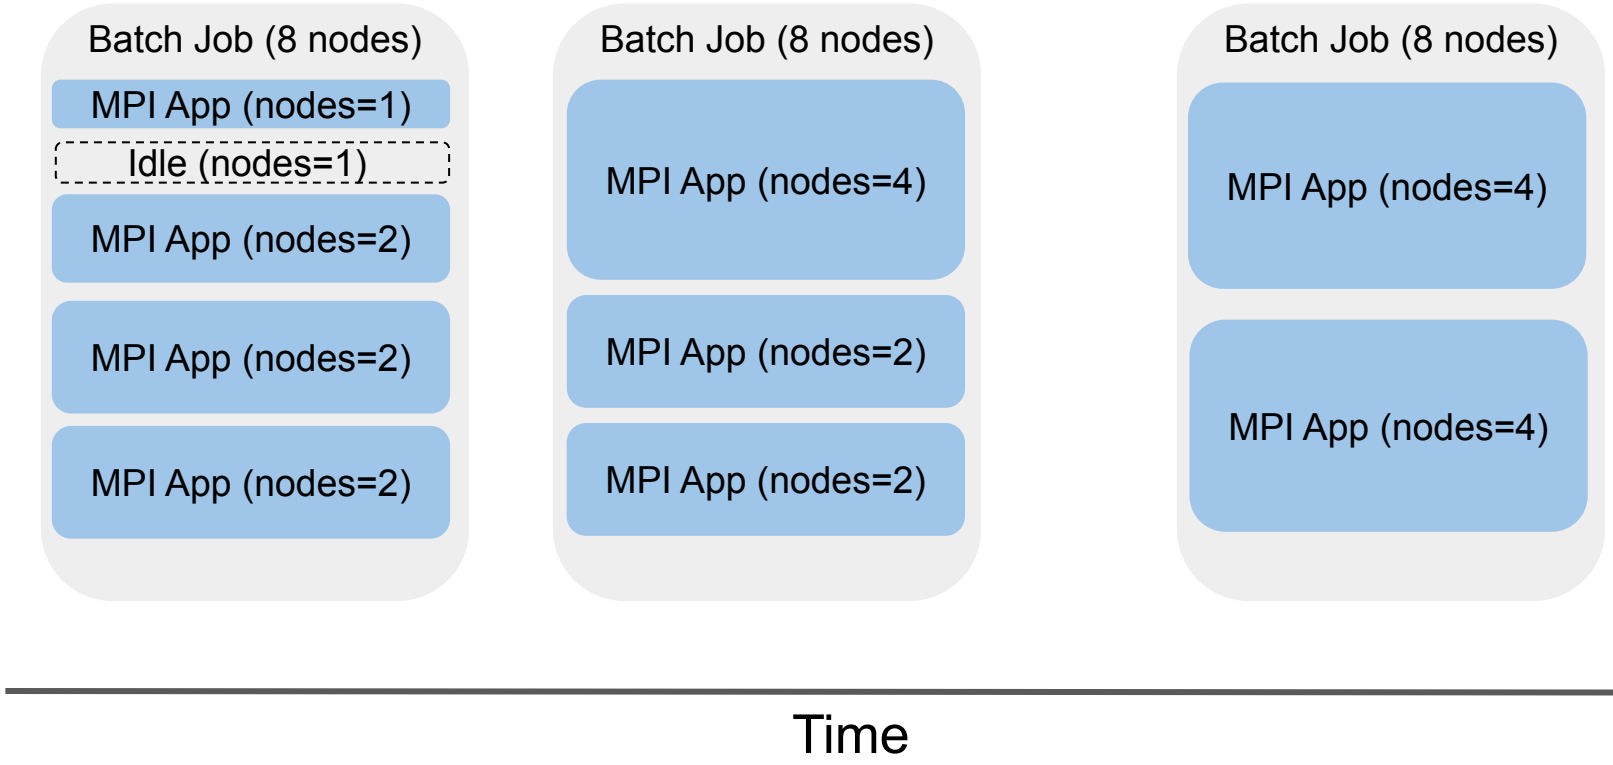

MPI in Parsl

mpi_executor.py

```
config = Config(
    executors=[
        MPIExecutor(
            address=address_by_interface('bond0'),
            max_workers_per_block=3,  # Assuming 2 nodes per task
            provider=PBSProProvider(
                account="parsl",
                worker_init=f"""module load miniconda; source activate /lus/eagle/projects/parsl/env""",
                walltime="1:00:00",
                queue="debug",
                scheduler_options="#PBS -l filesystems=home:eagle:grand",
                launcher=SimpleLauncher(),
                select_options="ngpus=4",
                nodes_per_block=6,
                max_blocks=1,
                cpus_per_node=64,
            ),
        ),
    ]
)
```

MPI in Parsl

mpiexecutor.py

```
config = Config(executors = [
    MPIExecutor(
        label="Polaris.MPIExecutor",
        # Specify the number of MPI apps that can run concurrently in a block
        max_workers_per_block=6,
        # Specify the mpi launcher type from (srun, aprun, mpiexec)
        mpi_launcher='mpiexec',
        provider=PBSProProvider(
            account="parsl",
            worker_init=f"""module load miniconda; source activate /lus/eagle/projects/parsl/env""",
            walltime="1:00:00",
            queue="debug",
            scheduler_options="#PBS -l filesystems=home:eagle:grand",
            launcher=SimpleLauncher(),
            select_options="ngpus=4",
            nodes_per_block=6,
            max_blocks=1,
            cpus_per_node=64,
        )
    )
])
```

MPI in Parsl

parsl_mpi_app.py

```
1  @bash_app
2  def lammps_mpi_application(infile: File, parsl_resource_specification: Dict):
3      # PARSL_MPI_PREFIX will resolve to `mpiexec -n 4 -ppn 2 -hosts NODE001,NODE002`
4      return f"$PARSL_MPI_PREFIX lmp_mpi -in {infile.filepath}"
5
6  # Resources in terms of nodes and how ranks are to be distributed are set on a per app
7  # basis via the resource_spec dictionary.
8  resource_spec = {
9      "num_nodes": 2,
10     "ranks_per_node": 2,
11     "num_ranks": 4,
12 }
13
14 future = lammps_mpi_application(File('in.file'), parsl_resource_specification=resource_spec)
```

MPI Support in Globus Compute: ShellFunctions

shell_function.py

```python
from globus_compute_sdk import ShellFunction, Executor

ep_id = "<SPECIFY_ENDPOINT_ID>"
# The cmd will be formatted with kwargs at invocation time
bf = ShellFunction("echo '{message}'")
with Executor(endpoint_id=ep_id) as ex:
    for msg in ("hello", "hola", "bonjour"):
        future = ex.submit(bf, message=msg)
        shell_result = future.result()  # ShellFunctions return ShellResults
        print(shell_result.stdout)
```

Ref: https://globus-compute.readthedocs.io/en/latest/executor.html#shell-functions

MPI Support in Globus Compute: MPIFunctions

mpi_function.py

```
executor.resource_specification = {
    'num_nodes': 2,      # Number of nodes required for the application instance
    'ranks_per_node': 4, # Number of ranks / application elements to be launched per node
    'num_ranks': 8,      # Number of ranks in total
}

mpi_func = MPIFunction("my_mpi_app")
future = executor.submit(MPIFunction("lammps -in lammps.in"))
```

Ref: https://globus-compute.readthedocs.io/en/latest/executor.html#shell-functions