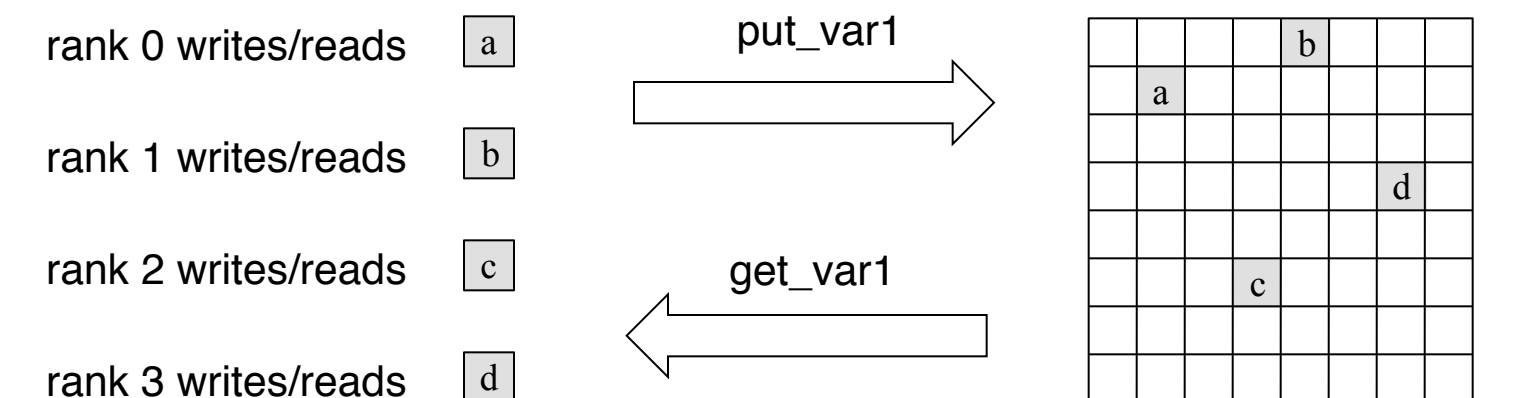
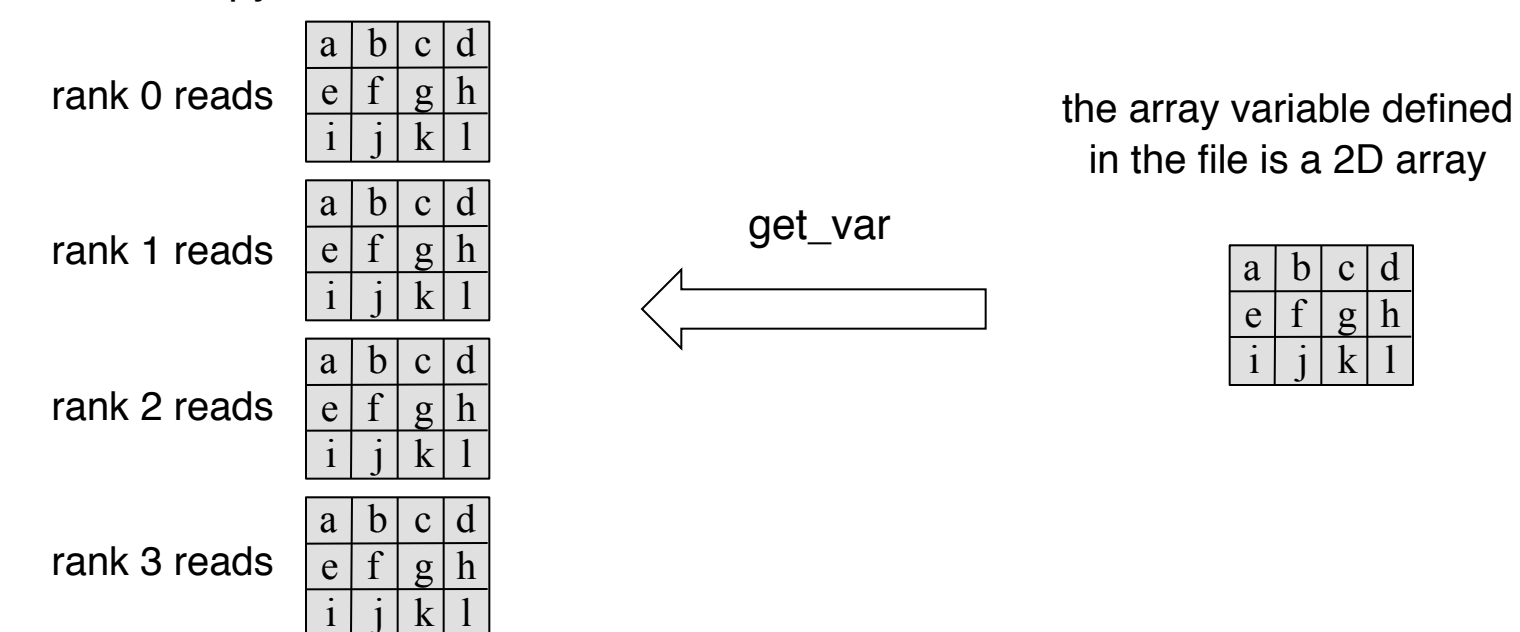


NetCDF APIs


each process writes 1 element at a different location



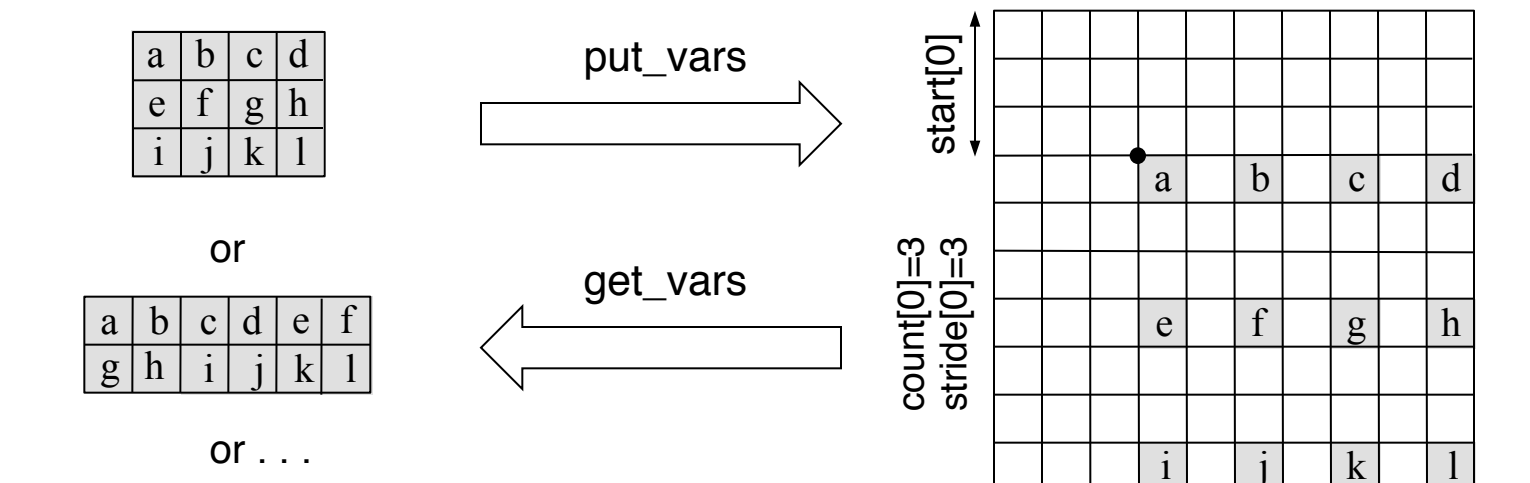
all processes get the same copy of a variable



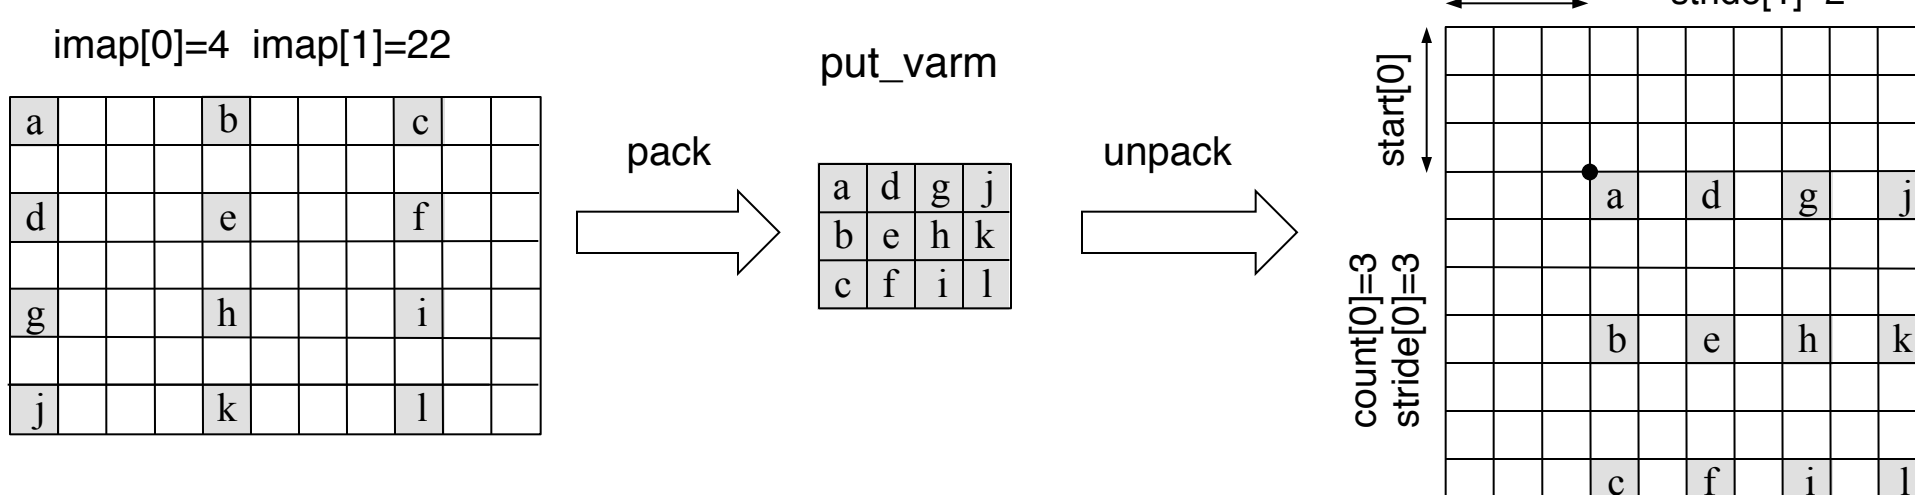
buffer in memory can be in any shape, but must be of size count[0]*count[1]



buffer in memory can be in any shape, but must be of size count[0]*count[1]



buffer in memory can be in any shape, in this example it is 7 x 11

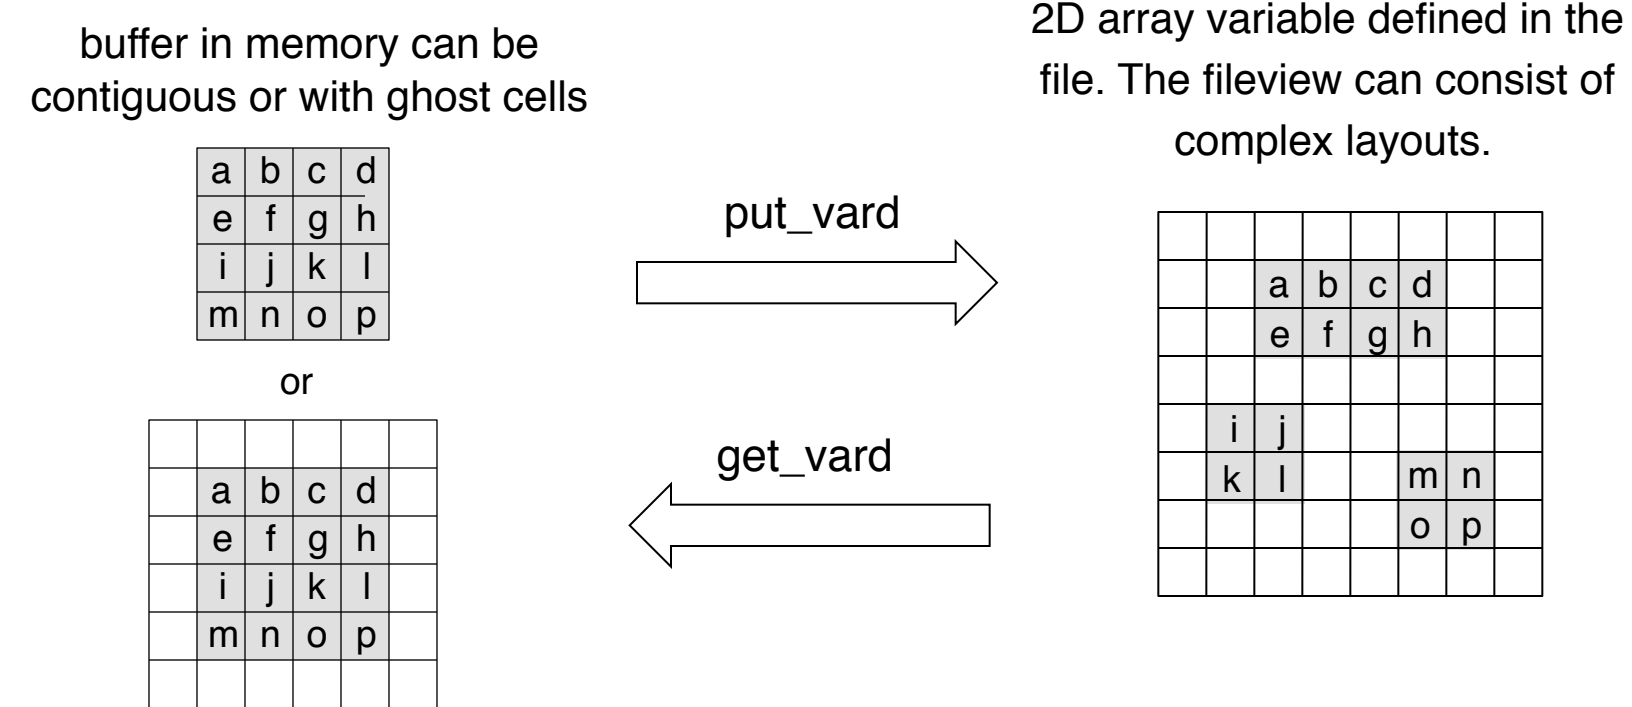
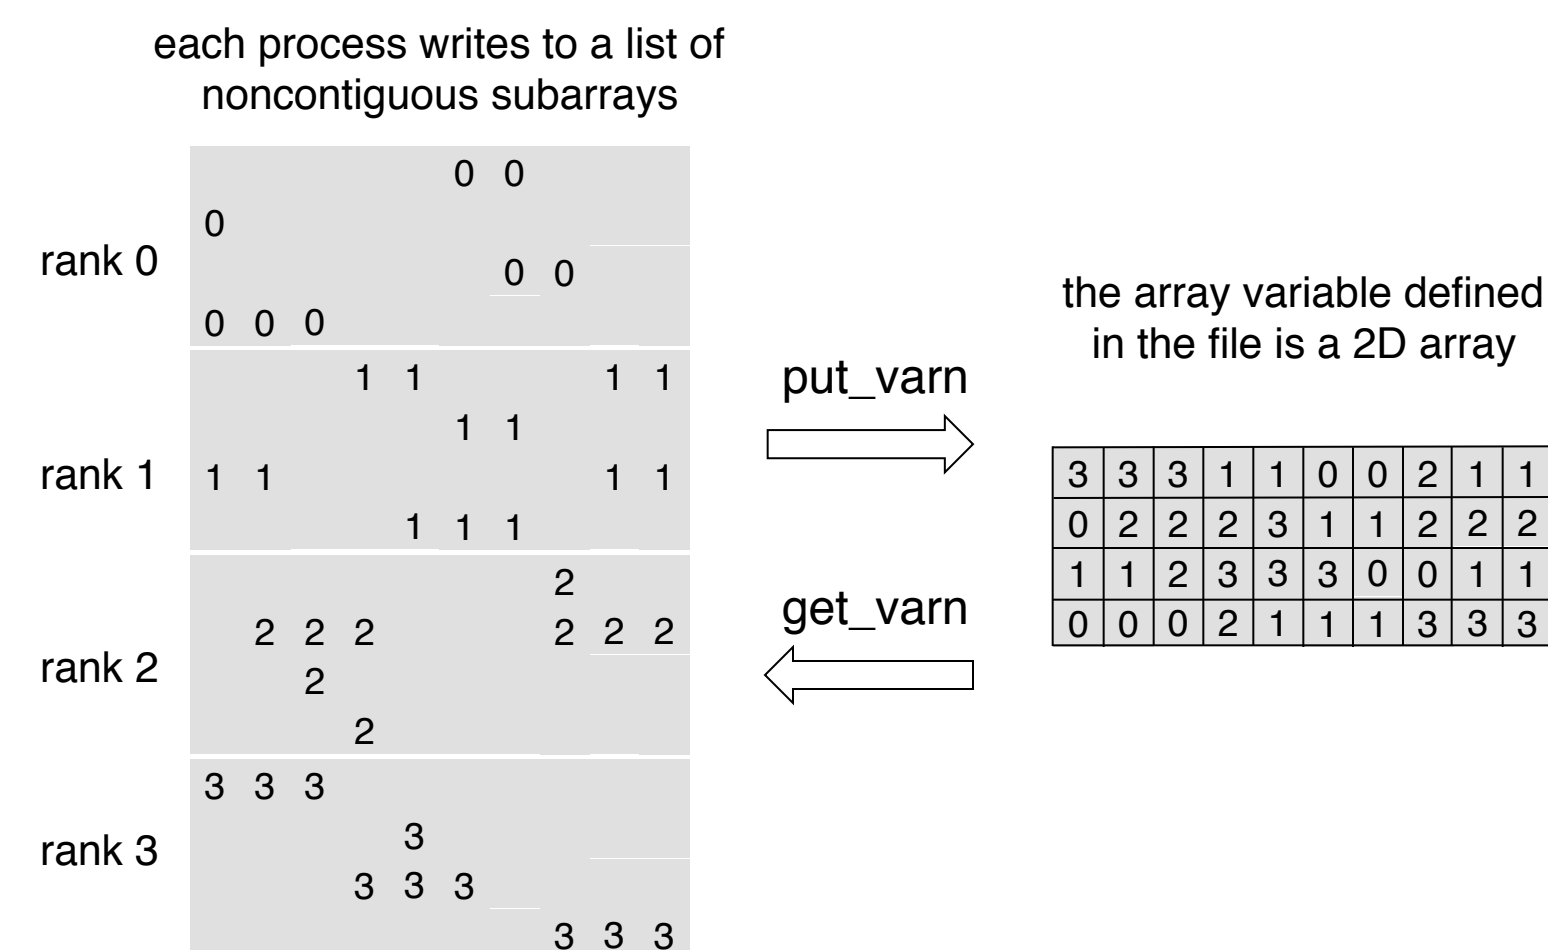


PnetCDF Nonblocking APIs


```

ncmpi_create          /* create netCDF dataset: enter define mode */
...
ncmpi_enddef         /* leave define mode */
ncmpi_iput_var<kind>_<type> /* post a nonblocking write request to a variable */
...
ncmpi_iput_var<kind>_<type> /* post another nonblocking request to the same or
...
ncmpi_wait_all       /* commit the posted nonblocking requests */
...
ncmpi_buffer_attach  /* tell PnetCDF the amount of space allowed for an internal write
...
ncmpi_bput_var<kind>_<type> /* post a buffered nonblocking write request to a variable */
...
ncmpi_bput_var<kind>_<type> /* post another buffered nonblocking write request to a variable */
...
ncmpi_wait_all       /* commit the posted nonblocking requests */
ncmpi_buffer_detach  /* tell PnetCDF to free the write cache */
ncmpi_close          /* close file */
    
```

PnetCDF varn and vard API Families



Parallel I/O with fill mode



- ❑ NetCDF fill mode is a mechanism to detect missing values
 - ✦ Variables may be written partially
 - ✦ Missing values can be user-defined or defaults defined in CDF
- ❑ Current implementation
 - ✦ Variables with fill mode enabled are first written with fill values, followed by true write requests from the application
 - ✦ For non-record variables, filling data is done at ncmpi_enddef()
 - ✦ For record variables, users are required to call ncmpi_fill_var_rec() explicitly to fill a record
- ❑ Performance impact
 - ✦ Causes a poor performance due to double writes
 - ✦ Relies on MPI-IO data sieving
 - ✦ Any solutions that try avoiding double writes can be challenging, because processes do not know requests from others, making detecting missing elements hard

PnetCDF used in Climate Applications

- ❑ CESM (Community Earth System Model)
 - ✦ NCAR/UCAR, sponsored by DOE and NSF
 - ✦ Reads/Writes CDF files through PIO library
- ❑ ACME (Accelerated Climate Modeling for Energy)
 - ✦ Joined work from 7 National Labs, sponsored by DOE
 - ✦ Reads/Writes CDF files through PIO library
- ❑ WRF (Weather Research and Forecasting) model
 - ✦ NCAR
 - ✦ Reads/Writes CDF files directly through PnetCDF
- ❑ PIO library
 - ✦ UCAR
 - ✦ Parallel I/O driver for CESM and ACME
 - ✦ User options to choose PnetCDF or netCDF-4
 - ✦ Captures various common I/O patterns from climate applications and translates them into PnetCDF function calls