Multiprocessor Coupler

Rationale:

- Speed increase

$$T_{comp\_C}(N_{proc\_C}) = T_{comp\_C}(1)/N_{proc\_C} + T_{ovrhd\_C}$$

("2D Coupler vs 3D components" justification for N_{proc_C}=1 only works with N_{proc} of the coupled system small enough compared to number of levels in the vertical)

- Multiprocess communication infrastructure relevant for Coupling-Without-a-Coupler project

- No changes in Components (so far as there is a Coupler)

Method:

- <u>tile per process</u>

+: full flexibility, with any N_{proc} operable

+: with only output fields tiled, no changes in the interpolation procedure

-: internal communication overhead

Requirements:

- Full backward compatibility
- Same results bitwise
- Still one process for communication with component models
- Tiling flexibility

Parallelized interpolation

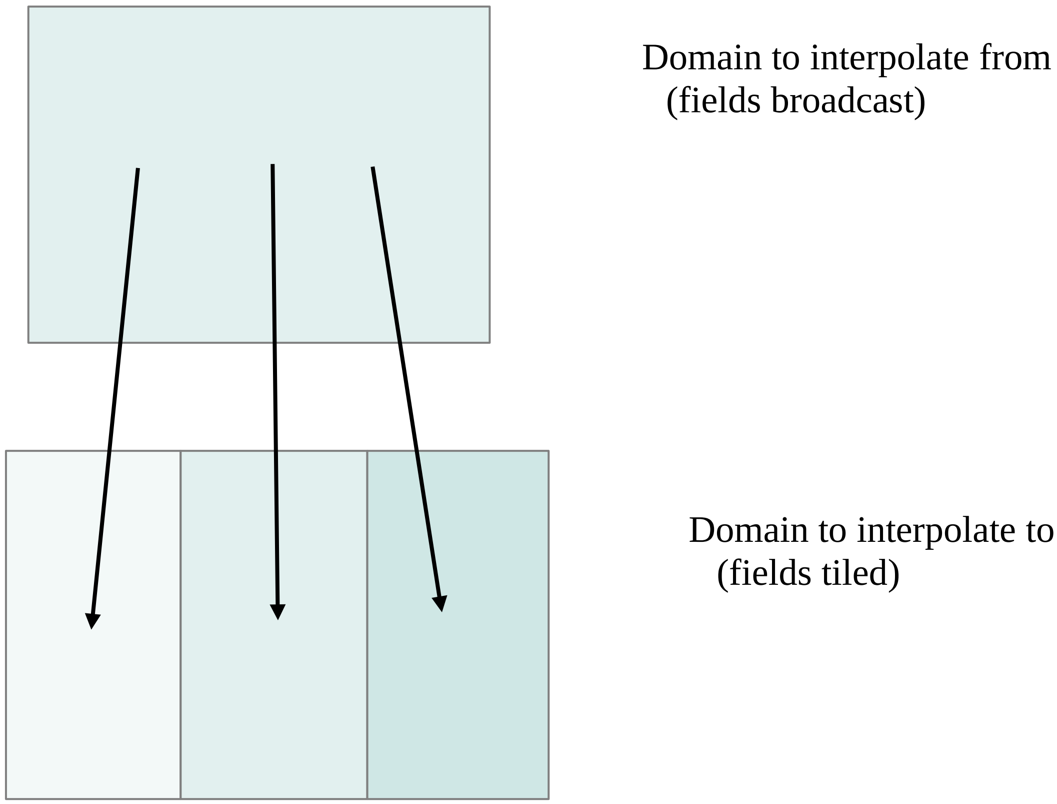

Method (same as in 1 process case):
– generalized bilinear interpolation (1^{st} order) if there are 4 valid sea point vertices in the grid cell to interpolate from;
– linear interpolation (1^{st} order) in the case of 3 valid sea point vertices;
– linear (0^{th} order) in the case of 2 valid sea point vertices (example below);
– constant (0^{th} order) in the case of 1 valid sea point vertex;
– no interpolation if no valid sea point vertices in the grid cell (background values or background values blended with extrapolation from other cells are used)\

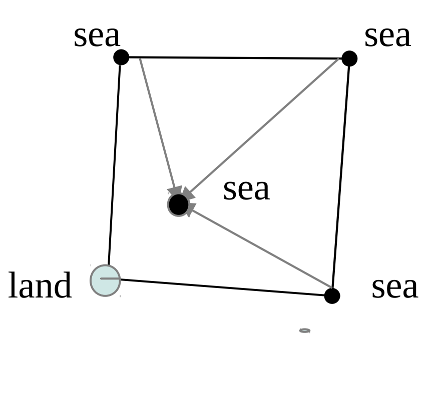

Extrapolation

Rationale: extrapolation is used in HWRF parent domain to provide smooth transition between (*) the SST interpolated from OM and used for HWRF gridpoints lying inside the OM domain, and (**) the background native HWRF SST used outside the OM domain.
It is also used to extrapolate surface fluxes to the areas in OM domain that lie outside the HWRF domain. This option is turned off in HYCOM, which uses GFS background surface fluxes instead; it may be harmlessly dispensed with in POM as well (or may be retained, as it is inexpensive). Only several grid rows adjacent to HWRF domain are processed.

Method. For each HWRF sea point outside OM domain, whose SST value is not yet determined: a vicinity is found that contains, in the same connected component of HWRF sea subdomain, such HWRF sea points that have already an SST value assigned via interpolation from OM domain or previous extrapolation; a mean SST value over these HWRF sea gridpoints is taken, and a weighted average of this value and the background HWRF native SST value is taken and assigned at the point at issue, the weight of the background SST being the greater the greater is the radius of the vicinity.

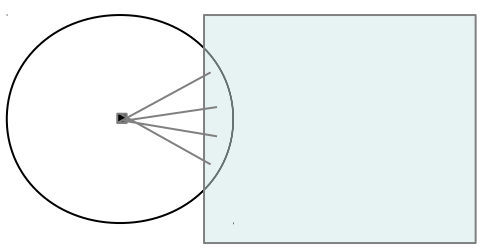

Issue: This extrapolation algorithm is not parallelizable. In the operational Coupler, it takes most of the computing time.

Solution: Extrapolation algorithm has been drastically accelerated in the multiprocessor Coupler, with the same results bitwise. (Another possible solution is to dispose of the algorithm; this, however, requires extensive testing to evaluate the impact on forecast skill scores.)

Speed (WCT), % OP

Time integration, HWRF+POM (dummies)

OP	100
1p, OP EXTRAP	100
1p, IMPROVED SST EXTRAP	24
2p, IMPROVED SST EXTRAP	22
2p, NO SST EXTRAP	23 (?)
3p, IMPROVED SST EXTRAP	21
2p, NO EXTRAP	18

Initialization of interpolation, HWRF+POM (dummies)

OP	190
1p = OP, no intp_init data storage	100
2p	47
3p	32
4p	24

Time integration, HWRF+POM+WW (dummies)

OP (=123% OP HWRF+POM)	100	
1p, IMPROVED SST EXTRAP	76	37
2p, IMPROVED SST EXTRAP	67	28

CONCLUSIONS / DISCUSSION

- The WallClockTime of the new 2-component (HWRF+POM) Coupler is some 25% or less of the operational Coupler WCT. (About the same acceleration is expected in the case of HWRF+HYCOM Coupler)

- Improved SST extrapolation in the new Coupler results in a drastic acceleration (some 75% reduction of time loop WCT), about the same as no SST extrapolation at all. Dispensing with SF extrapolation would result in further but minor WCT reduction (3% to 4% of the operational WCT)

- Parallelization of the time loop (interpolation and analysis) turns out poorly scalable, the total time loop WCT in the new Coupler with 2 processes being some 22% of the operational WCT (i. e. additional 3% acceleration is gained). The probable reason is that interpolation is fast enough (or MPI broadcast is slow enough) for interpolation parallelization gains to be offset by MPI broadcast losses

- Parallelization of the (pre-time-loop) initialization of interpolation turns out ideally scalable, WCT being equal to op. WCT divided by the number of processes

- Results on time loop acceleration for the new Coupler in the case of 3 components (HWRF_POM+WW) are ambiguous, the WCT gain randomly jumping from some 25 to 75%. A reduction and stabilization of WCT may be hoped for in the case of real component models, due to more opportunities to MPI-communicate during computations. However, random WCT variability and/or losses on MPI waiting are probably inherent to “hub coupling” in the case of more than 2 component models (which is another argument in favor of “Coupling without a Coupler”)

- To attempt improvement of time loop scalability and possibly performance in general, a switch to nonblocking MPI sends is advisable (if not in the current Coupler then for “couplerless coupling”)