ESMF Component Overhead in CCSM 4.0

Peggy Li
May 12, 2010

Objectives

In the latest release of the Community Climate System Model (CCSM 4.0, released on April 1, 2010), a new ESMF compliant component interface was added in addition to the MCT-based CCSM designed component interface. The ESMF interface implementation exists in CCSM to support further development and testing of an ESMF driver or ESMF couplers and to allow CCSM model components to interact with other coupled systems using ESMF coupling standards.

The objective of this benchmark is to determine the overhead of the ESMF components in a selected set of standard CCSM test cases and data resolutions. The performance goal is to keep the ESMF component overhead within 3% of the total run time.

The benchmark was conducted on the Cray XT4, jaguar, at Oak Ridge National Lab and the IBM Power 575, bluefire, at UCAR. The software versions are CCSM_4_0_0_beta42 and ESMF_5_0_0_beta_snapshot_01. The CCSM compset used for the benchmark is B1850CN and the data resolution is f09_g16. B1850CN includes the atmospheric model CAM, the land model CLM, the ice model CICE, the ocean model POP2 and a coupler. The data resolution is 1.25 degree x 0.9 degree global grid with 17 vertical levels for both the atmospheric and land model, i.e. 288x192x17 grid. The data resolution for the ocean model is 320x384x60.

Benchmark Programs

Benchmark results are shown for only one CCSM4 configuration, B1850CN, although three configurations were tried. The reason is that these benchmarks were performed while CCSM4 was still in development, and several configurations did not run at that time with ESMF or MCT or either.

We used CCSM script *create_newcase* to build the test case. For example, to build a test case of B20TRCN with ESMF interface and f09_g16 data resolution on jaguar using 64 PEs, the script looks like the following:

```
create_newcase -mach jaguar -res f09_g16 -compset B20TRCN
-pes_file pes64.xml -case B20TRCN.ESMF.64
cd B20TRCN.ESMF.64
xmlchange -file env_run.xml -id STOP_N -val 20
xmlchange -file env_run.xml -id REST_OPTION -val never
xmlchange -file env_build.xml -id USE_ESMF_LIB -val TRUE
xmlchange -file env_build.xml -id COMP_INTERFACE -val ESMF
```

```
configure -case
B20TRCN.ESMF.64.jaguar.build
```

Where *pes64.xml* is supplied by the user to specify how the data is decomposed in each component. To build with the ESMF component interface, we need to modify the build environment file *env_build.xml* to instruct CCSM to use the ESMF library. We also changed *env_run.xml* to run the model for 20 days without using a restart file.

The Result

The results reported here are based on the B1850CN compset with data resolution f09_g16.

We ran the benchmark on jaguar and bluefire using 64, 128, 256 and 512 PEs. We used the timing results reported by the timing summary files generated by the test run. Table 1 is the initialization time, the total model run time and the run times for each component (the land component (lnd), ice component (ice), atmospheric component (atm), ocean component (ocn) and the coupler (cpl)) from both jaguar and bluefire on 4 different configurations for a 20 day simulation. The blue colored columns are the timing difference in percentage between the CCSM/ESMF and CCSM/MCT. The yellow highlighted rows are the total run time. Note the overall overhead for the ESMF component interface is between 0.09% to 3.13%. There is only one configuration (512 PE bluefire) that exceeds the 3% overhead requirement.

Figure 1 and Figure 2 compares the init time and the total run time from the CCSM/MCT run and the CCSM/ESMF runs on bluefire and jaguar, respectively. The ESMF and MCT curves for the total run time are almost entirely overlapped because the difference is very small.

Conclusion

In summary, the run time differences between MCT and ESMF are very small on both machines and on all the configurations. They are all within 3% except the 512PE bluefire run., i.e., 3.12%. The init time differences between MCT and ESMF are bigger. It looks like that ESMF init time increases more rapidly with the increase of the processors, therefore, the ESMF init time gets worse when large number of processors are used (>=256).

		Bluefire			Jaguar		
# proce	models	ESMF	MCT	difference(%)	ESMF	MCT	difference(%)
64	init	71.042	66.626	6.628043	46.961	46.884	0.164235
	total run	1754.933	1744.835	0.578737	2317.16	2261.379	2.466681
	lnd	105.869	94.636	11.86969	153.06	129.733	17.98078
	ice	179.464	176.947	1.42246	208.106	206.718	0.671446
	atm	1141.129	1140.786	0.030067	1352.081	1331.204	1.56828
	ocn	308.222	308.871	-0.21012	555.926	554.636	0.232585
	cpl	18.837	18.303	2.917554	43.349	43.531	-0.418093
# proce	models	ESMF	MCT	difference(%	ESMF	MCT	difference(%)
128	init	69.947	70.943	-1.403944	44.995	46.023	-2.233666
	total run	981.246	962.617	1.935245	1305.625	1282.293	1.819553
	lnd	59.598	50.992	16.87716	78.609	73.235	7.338021
	ice	105.189	102.62	2.503411	133.138	131.799	1.015941
	atm	631.657	627.107	0.725554	751.4	744.455	0.932897
	ocn	171.576	170.517	0.621052	307.287	306.747	0.176041
	cpl	11.909	11.451	3.999651	26.014	24.541	6.0022
# proce	models	ESMF	MCT	difference(%	ESMF	MCT	difference(%)
256	init	77.25	75.09	2.876548	50.834	51.142	-0.602245
	total run	573.15	560.724	2.216064	755.784	755.09	0.09191
	lnd	98.766	93.026	6.170318	114.363	109.529	4.413443
	ice	81.584	78.546	3.867797	92.447	90.12	2.582113
	atm	343.255	339.005	1.253669	404.971	405.535	-0.139076
	ocn	104.49	104.201	0.277349	191.662	195.367	-1.896431
	cpl	9.754	9.276	5.153083	19.529	19.232	1.544301
# proce	models	ESMF	MCT	difference(%	ESMF	MCT	difference(%)
512	init	86.955	76.411	13.79906	54.163	52.697	2.781942
	total run	372.81	361.496	3.129772	516.132	509.088	1.383651
	lnd	69.983	65.145	7.42651	80.882	77.045	4.980206
	ice	61.363	58.424	5.030467	71.362	69.865	2.142704
	atm	209.299	205.271	1.962284	256.472	256.22	0.098353
	ocn	70.917	70.885	0.045144	141.287	141.506	-0.154764
	cpl	8.74	7.974	9.60622	17.232	17.098	0.783717

Table 1: CCSM Timing comparison Compset B1850CN, f09_g16 (all the times are in seconds)

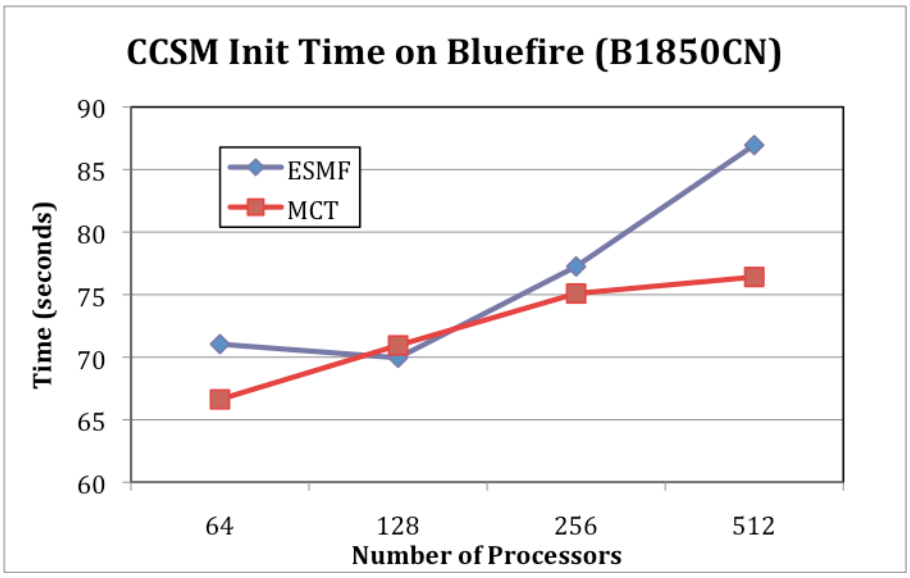

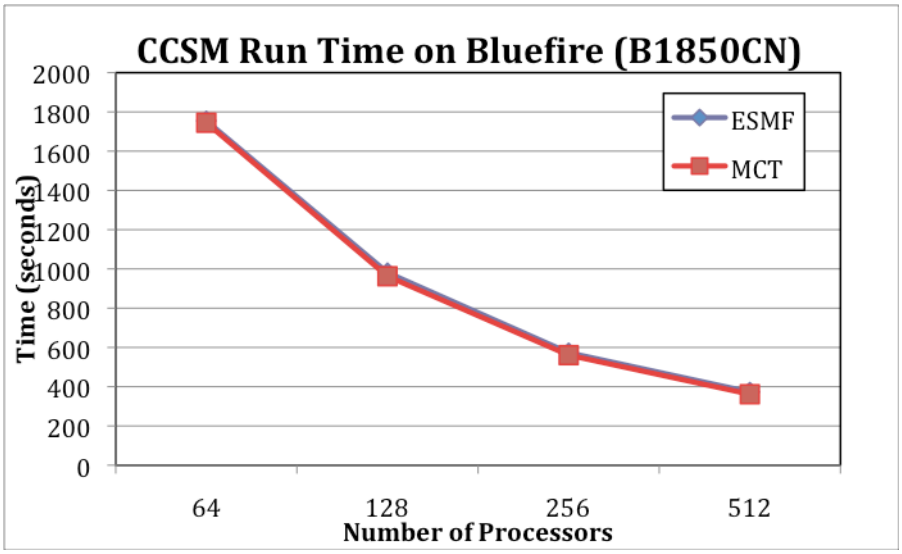

Figure 1 The CCSM Init Time and Total Run Time on bluefire

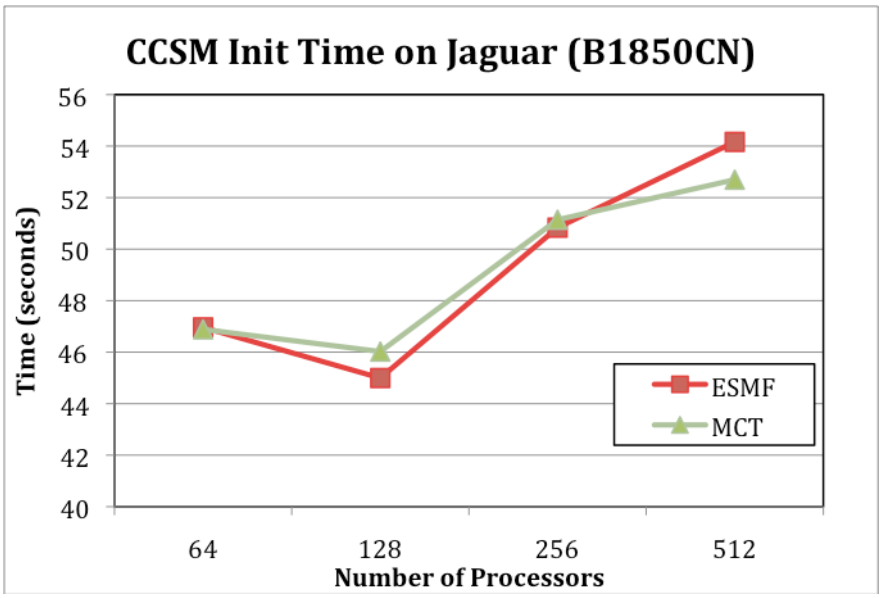

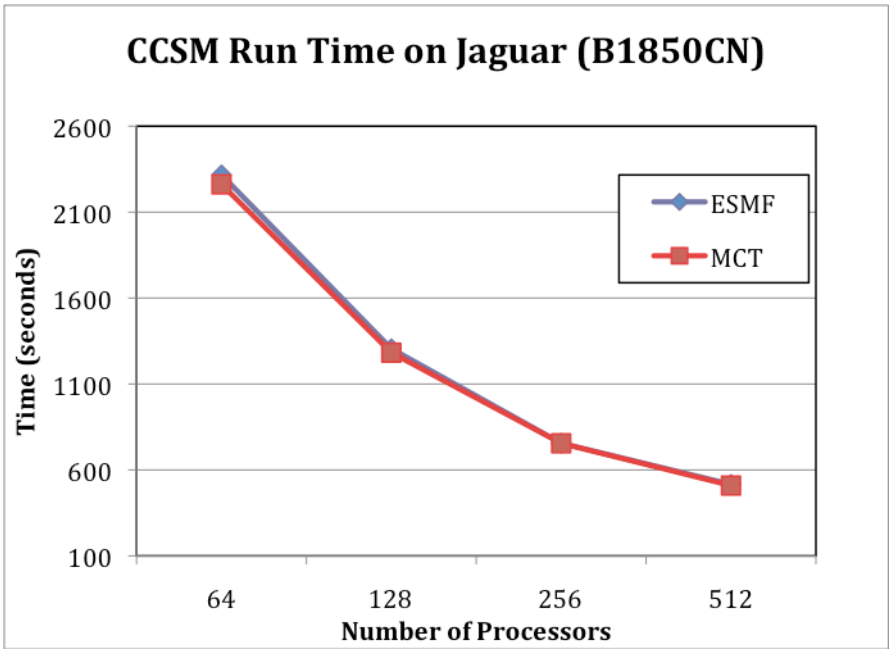

Figure 2 The CCSM Init and the total run time on jaguar