

Performance Analysis Exercises with Vampir



Exercise Trace Files

% ls /home/nct00/nct00006/trace-examples

```
01-p100-cosmo-specs-orig
02-p100-cosmo-specs-fd4
03_wrf_deimos
04_sbmfd4_jugene
```

- Four trace files for exercising performance analysis with Vampir
- Traces show real application runs
- Do the traces contain performance problems?
- If yes, try to find their causes

	3	2			5 🗟	8	i 🎽	5	3 🚯	*	Q 0 s				5	78.101 s			578 s
							Timelir	ne						××			Function Summ	ary	
		0 s 5	50 s	100 s	150 s	200 s	250 s	300 s	350 s	400 s	s 450 s	500 s	550 s		All Processes, Acc	cumulated	Exclusive Time	per Function G	roup
Proces	s 0	ODDI	mm											×	20,000 s	15,000 s	10,000 s	5,000 s	
Proces	s 1	00000	DDD											×	18 78	3 741 s			MP
Proces	s 2	ÖDDD	$\overline{\mathrm{mm}}$			nnnn		000000		20000					10,70	0.7410	11.577.183	5	MP UTIL
Proces	s 3	ODDI	nn			nnin					00000000		0000000					1,550.646 s	Application
Proces	s 4	0000	DDD	\mathbf{m}	$\overline{\mathbf{m}}$	nnin	min			∞			0000000					1,290.593 s	COUPLE
Process	s 5	ODDI	$\overline{\mathrm{mm}}$		$\overline{\mathbf{m}}$	\mathbf{p}			$\infty \infty$	∞		\mathbf{m}						993.146 s	METEO
Proces	s 6	ODD	mm	DDD	$\overline{m}\overline{m}$	\mathbf{m}				∞		$\infty \infty \infty$						4.366	S s VT_API
Proces	s 7	00001	$\mathbf{D}\mathbf{D}\mathbf{D}$	$\mathbf{D}\mathbf{D}\mathbf{D}$		DDDD				∞									
Process	s 8	\mathbf{ODD}	$\overline{\mathrm{DDD}}$	\mathbf{n}	DDDDD	ntint			0000000	∞									
Proces	s 9	00000	\mathbf{m}	DDDD	mmm	ppp	\mathbf{m}	mm		∞							Context View		
Proces	s 10	00001	$\mathbf{D}\mathbf{D}\mathbf{D}$	\mathbf{m}	$\infty \infty$					∞		$\infty \infty \infty \alpha$			-	×	Trace		× +
Process	s 11	03331	0001	9 000				0000000		$\mathcal{D}\mathcal{D}\mathcal{D}$			1000000		Property		Value		
Process	s 12	00000	\overline{nn}							$\mathcal{D}\mathcal{D}\mathcal{D}\mathcal{D}$					File		/Users/maweb	er/ZIH/Traces/V	ampir-Tutorial
Proces	s 13	000	000	0000	mm	0000	000000				8080808				Creator		VampirTrace F	: 11 2	
Proces	s 14	0000			mm	nnn					8080808								
Proces	s 15					$\mu\mu\mu$,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,								version		1.11 alpha		
Proces	s 16	0,777						$\frac{1}{1}$			uuuu maana				Number of Proces	sses	100		
Proces	s 17														Timer Resolution		50 ns		
Proces	s 18			<u> </u>											Trace Times				
Proces	s 19		****	****											Start		Thu Nov 312	36:57 2011 (132	2032021776169
Process	s 20														Stop		Thu Nov 312	46:30 2011 (132	203207909206
Process	s ∠1 ∈ 22														Elapsed		00:09:33 (573	158989)	
Proces	s 22	000	mb			nnn	nnn	mm							VampirTrace Envi	ironment	,		
Proces	s 24	om	m	m		nnn	nnn	mm							VT MODE		TRACE		
Process	s 25	0000	min	mm		nnn										76	20014		
Proces	s 26	00000	nini	nnn											VI_BUFFER_SI		2001/0		
Proces	s 27	ODDI	nm	2000	CCCCCCC										VT_SYNC_FLUS	SH	no		
Proces	s 28	ODDI	mm		CCCCCCCC										VT_SYNC_FLUS	SH_LEVEL	80		
Proces	s 29	ODDI		20000	CCCCCCCC										VT_MAX_FLUS	HES	1		
Proces	s 30	00000	$\overline{\mathbf{m}}$			nin			0000000	20000					VT_METRICS		PAPI_FP_OPS		
															VT_RUSAGE		<not set=""></not>		

- Weather forecast code
 COSMO-SPECS
- Run with 100 processes
- COSMO: weather model (METEO group)
- SPECS: microphysics for accurate cloud calculation (MP and MP_UTIL group)
- Coupling of both models done in COUPLE group



- Compared to METEO, MP and MP_UTIL are very compute intensive, however this is due to more complex calculations and no performance issue
- Problem: >32% of time spent in MPI
- MPI runtime share increases throughout the application run



- Zoom into the first three iterations
- MP/MP_UTIL perform four sub-steps in one iteration
- Low MPI time share
- Everything is balanced and looks okay



- Zoom into the last three iterations
- Very high MPI time share
 (>50%)
- Large load imbalance caused by MP functions around Process 54 and Process 64



- **PAPI_FP_OPS** counter showing higher FLOPs rates on processes causing the imbalance
- Reason for imbalance: Static grid used for distribution of processes.
 Depending on the weather, expensive cloud computations (MP group) may be only necessary on some processes

02-p100-cosmo-specs-fd4

	🐹 🔛	i i i	3 🔢 👪 🗶) 🐴 🐚 🕻	۹ 🖈 象	S		276.026 s		276 s
	_		Time	line				Function	n Summary	
	0 s	50 s	100 s	150 s	200 s	250 s	All Processes	, Accumulated Exclusi	ve Time per Functi	ion Group
Process 0							19,207,026,0	0 s 10,000 s	5,000 s	US
Process 1							18,307.920 \$		3.0	865.817 S MD LIT
Process 2			In the second		000000000000000000000000000000000000000				3.457.412 s	MP_01
Process 3		וכתכתכתכו	CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC	COCCECCECCECCE	DDDDDDDDDDDDDDD	CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC			1,443.0	047 s COUPL
Process 4	• • • • • • • • • • • • • • • • • • • •		COCOCOCOCOCOCOCOCOCOCOCOCOCOCOCOCOCOCO	CCCCCCCCCCCCCCCC	DDDDDDDDDDDDDD				3	93.818 s 📕 METEO
Process 5										114.399 s 🛛 FD4
Process 6					COLOCOCOCOCOCOCOCO	COCOCOCOCOCOCOCOCOCOCO				3.484 s VT_API
Process 7			COCOCOCOCOCOCOCOCOCOCOCOCOCOCOCOCOCOCO	ממככמתמכמים	CONCERCENCE	CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC				
Process 8		וככתכככככ	ICCOCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC	acceleter and	BODDODDDDDDDD D	CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC				
Process 9	• • • • • • • • • •	ICCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC	CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC	COCCCCCCCCCCCCC	DDDDDDDDDDDDDDDD					
Process 10										
Process 11										
Process 12			ICCOLOCOCOCOCOCOCOCOCOCOCOCOCOCOCOCOCOCO		000000000000000000	000000000000000000000000000000000000000				
Process 13			000000000000000000000000000000000000000		000000000000000000	000000000000000000000000000000000000000				
Process 14			פטרמפרמכומ	תוכנותותות			•			
Process 15			ת היה היה היה היה היה היה היה היה		סוונוווווווווווווווווווו					
Process 16			ענייייייייייייייייייייייייייייייייייייי	\dots	\mathbf{u}			-	•	
Process 17			,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,			עררערררעררר		Conte	ext View	
Process 18							-			· · ·
Process 19							Property		Value	
Process 20							File		/Users/maw	veber/ZIH/Traces/V
Process 21							Creator		VampirTrac	e 5.13b111102
Process 22							Version		1.11 alpha	
Process 23							Number of Pr	ocesses	100	
Process 24							Timor Posolu	tion	50 ps	
Process 26	• • • • • • • • • • • • • • • • • • • •					mmmmm	Trace Times		50 113	
Process 27						and a second	Trace Times		-	10.00.10.0011 (100
Process 28							Start		Mon Nov 7	10:38:19 2011 (132
Process 29						,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,	Stop		Mon Nov 7	10:42:55 2011 (132
Process 30		\dots					Elapsed		00:04:36 (2	276109382)
							VT_METRIC	cs	PAPI_FP_OF	PS
								-		

- Weather forecast code COSMO-SPECS
- Run with 100 processes
- COSMO: weather model (METEO group)
- SPECS: microphysics for accurate cloud calculation (MP and MP_UTIL group)
- Coupling of both models done in COUPLE group
- Dynamic load balancing (FD4 group)

02-p100-cosmo-specs-fd4

: 🗮 🔝		🛄 🐻 🔜 🍮 🔁) 🔠 🏬 🔺	é 🐴 🐚	🔯 🖈 💡 🔤			276.0	26 s		276 s
			Tim	eline			××		Function Sum	mary	
	0 s	50 s	100 s	150 s	200 s	250 s		All Processes, Accumu	lated Exclusive Tin	ne per Function Gro	up
Droopen 0	-			•			1	15,000 s	10,000 s	5,000 s	0 s
Process 0							×	18.307.926 s (66.37%)		MP
Process 6	_						~		3 865 817 c (1	4 01%)	
Process 9	-1 1								2 457 412 - (4.0178)	
Process 12	-								3,457.412 S (12.53%)	MPI
Process 15	-								1,443	.047 s (5.23%)	COUPLE
Process 18	-1 1									393.818 s (1.43%)	METEO
Process 21										114.399 s (0.41%)	FD4
Process 24										3 484 \$ (0 01%)	
Process 27										0.404 0 (0.0170)	
Process 33											
Process 36	-										
Process 39	-1 1										
Process 42	-										
Process 45	-										
Process 48	-				7						
Process 51					2 2 2 2 2 2 2 2 2 2 2 2 2 2 2 2 2 2 2 2						
Process 54											
Process 57											
Process 63								0			
Process 66	-										
Process 70	- 1										
Process 73	-								Context Vie	W	
Process 76	-								× Trace		× +
Process 79	-										
Process 82	-]							Property		Value	
Process 85								File		/Users/maweber/7	IH/Traces/Vai
Process 88										,	,
Process 91								Creator		VampirTrace 5.13b	5111102
Process 94 Process 97	-							Version		1.11 alpha	
All Processes	: Exclus	: ive Time per Function Grour	,	•				Number of Processes		100	
								Timer Resolution		50 ns	
100	%							Trace Times			
60	%							Start		Mon Nov 7 10:38:	19 2011 (1320
40	%			MP				Stop		Mon Nov 7 10:42:	55 2011 (1320
20	%							Flanced		00.04.26 (276400)	202)
20	%							ciapsed		00.04.30 (2/0109)	302)
0	/0							VT_METRICS		PAPI_FP_OPS	
69.5 s											

- Dynamic load balancing mitigates the balance problems of the original COSMO-SPECS version
- MPI time share is reduced to <13%
- MPI time share stays constant throughout the application runtime
- Runtime reduced by factor of 2.1, from initially 578s to 276s

02-p100-cosmo-specs-fd4



- Zoom into last three iterations
- FD4 balances MP load (precipitation processes in clouds) across all available processes

03_wrf_deimos

Understand Function Summary Process 0 50% 20% 30% 40% 50% 60% Process 1 50% 60% 750% 0 750% 0 Process 1 50% 60% 750% 0 0 750% 0 0 Process 3 50% 60% 750% 0			2	G		6		E 88 .	ا 🛃	§ 19		*	Q	0 s	r		n i an	64	4.896 s	In haut I			65 s
0 s 10 s 20 s 30 s 40 s 50 s 60 s Process 0 10 s 20 s 30 s 40 s 50 s 60 s Process 1 10 s 10 s 10 s 0 s 0 s 10 s 0 s Process 2 50 s 10 s 10 s 10 s 0 s 0 s 10 s 0 s Process 2 50 s 10 s 10 s 10 s 10 s 0 s 10 s								Timeline											Function	Summary			
Process 0 mm		0	S	10 s			20 s	30 s		40 s		50 s		60 s		All Process	es, Accumi	lated Exc	lusive Time	e per Funct	tion Group		
Process 0	Dresses	0	ev ne			YAAY			1111		-					1,5	00 s	75	50 s	0	s		
Process 1 010 011	Process	1	sync													1,904.449	S				DYN		
Process 2 Ordel Ordel Process 3 VTR Process 1 VTR Process 5 A Treeco	Process	ו ס	Sync														1,294	4.59 s	074.004		MPI		
Process 3 9702 223.953 g VI_API Process 5 9702 0	Process	2	Sync																671.993	3 S	PHYS		
Process 6 Sync Context View Process 7 Sync Context View	Process	3 1	sync																228.95	03 S	VT_API		
Process 3 Sync Context View Process 6 Sync Context View Process 10 Sync Context View Process 12 Sync Context View Process 13 Sync Context View Process 14 Sync Context View Process 15 Sync Context View Process 16 Sync Context View Process 13 Sync Context View Process 14 Sync Context View Process 13 Sync Context View Process 14 Sync Context View Process 12 Sync Context View Process 23 Sync Context View Process 24 Sync Context View<	Drocess	4 E	sync																	44.95 5			
Process 0 9mc Context View Process 1 9mc Context View Process 2 9mc Context View Process 2 9mc Context View Process 2 9mc Context View Process 2 <td>Process</td> <td>5</td> <td>Sync</td> <td></td> <td>1.230 \$</td> <td>IO_INETCI</td> <td></td> <td></td>	Process	5	Sync																	1.230 \$	IO_INETCI		
Process 0 Sync Context View Process 10 Sync Context View Process 11 Sync Context View Process 12 Sync Context View Process 13 Sync Context View Process 14 Sync Context View Process 15 Sync Context View Process 16 Sync Context View Process 17 Sync Context View Process 19 Sync Context View Process 10 Sync Context View Process 10 Sync Context View Process 17 Sync Context View Process 20 Sync Context View Process 21 Sync Context View Process 22 Sync Context View Process 23 Sync Context View Process 24 Sync Context View	Process	0 7	sync																				
Process 0 ymc Context View Process 10 ymc Context View Process 11 ymc Context View Process 12 ymc Context View Process 13 ymc Context View Process 14 ymc Context View Process 15 ymc Context View Process 14 ymc Context View Process 15 ymc Context View Process 16 ymc Context View Process 17 ymc Context View Process 18 ymc Context View Process 17 ymc Context View Process 18 ymc Context View Process 19 ymc Context View Process 20 ymc Context View Process 21 ymc Context View Process 22 ymc Context View Process 23 ymc Context View Process 24 ymc Context View Process 25 ymc Context View Process 26 ymc Context View	Process	/ 0	sync																				
Process 10 Sync Context View Process 11 Sync Context View Process 12 Sync Context View Process 13 Sync Context View Process 14 Sync Context View Process 15 Sync Context View Process 16 Sync Context View Process 16 Sync Context View Process 17 Sync Context View Process 18 Sync Context View Process 16 Sync Context View Process 17 Sync Context View Process 18 Sync Context View Process 18 Sync Context View Process 19 Sync Context View Process 21 Sync Context View Process 22 Sync Context View Process 23 Sync Context View Process 24 Sync Context View Process 25 Sync Context View Process 24 Sync Context View Process 25 Sync Context View <td>Drocess</td> <td>0</td> <td>Sync</td> <td></td>	Drocess	0	Sync																				
Process 10 Sync Context View Process 12 Sync Context View Process 13 Sync Context View Process 14 Sync Context View Process 15 Sync Context View Process 16 Sync Context View Process 17 Sync Context View Process 18 Sync Context View Process 17 Sync Context View Process 18 Sync Context View Process 19 Sync Context View Process 10 Sync Context View Process 17 Sync Context View Process 20 Sync Context View Process 21 Sync Context View Process 22 Sync Context View Process 23 Sync Context View Process 24 Sync Context View Process 25 Sync Context View Process 26 Sync Context View Process 27 Sync Context View Process 28 Sync Context View <td>Process</td> <td>9 10</td> <td>sync</td> <td></td> <td>XX</td> <td></td>	Process	9 10	sync		XX																		
Process 11 Sync Context rew Process 13 Sync Context rew Process 14 Sync Context rew Process 15 Sync Context rew Process 16 Sync Context rew Process 17 Sync Context rew Process 18 Sync Context rew Process 17 Sync Context rew Process 18 Sync Context rew Process 17 Sync Context rew Process 18 Sync Context rew Process 17 Sync Context rew Process 18 Sync Context rew Process 17 Sync Context rew Process 17 Sync Context rew Process 18 Sync Context rew Process 20 Sync Context rew Process 21 Sync Context rew Process 23 Sync Context rew Process 24 Sync Context rew Process 25 Sync Context rew Process 26 Sync Context rew	Process	10	sync		ă														Conte	● xt View			
Process 12 Sync Composition of the sync o	Process	12	sync		in the				1111			n n				-		×					7 +
Process 10 symc O <	Process	12	sync		n n				$\frac{1}{1}$				γ			-							
Process 15 Sync Cost State //Jeers/maweber/ZIH/Traces/Vampir-Tutorial-Analysis-Exampler Process 16 Sync Cost State //Jeers/maweber/ZIH/Traces/Vampir-Tutorial-Analysis-Exampler Process 16 Sync Cost State //Jeers/maweber/ZIH/Traces/Vampir-Tutorial-Analysis-Exampler Process 17 Sync Cost State //Jeers/maweber/ZIH/Traces/Vampir-Tutorial-Analysis-Exampler Process 18 Sync Cost State //Jeers/maweber/ZIH/Traces/Vampir-Tutorial-Analysis-Exampler Process 19 Sync Cost State //Jeers/maweber/ZIH/Traces/Vampir-Tutorial-Analysis-Exampler Process 19 Sync Cost State //Jeers/maweber/ZIH/Traces/Vampir-Tutorial-Analysis-Exampler Process 20 Sync Cost State //Jeers/maweber/ZIH/Traces/Vampir-Tutorial-Analysis-Exampler Process 21 Sync Cost State //Jeers/maweber/ZIH/Traces/Vampir-Tutorial-Analysis-Exampler Process 22 Sync Cost State //Jeers/maweber/ZIH/Traces/Vampir-Tutorial-Analysis-Exampler Process 25 Sync Cost State //Jeers/maweber/ZIH/Traces/Vampir-Tutorial-Analysis-Exampler Process 25 Sync Cost State //Jeers/maweber/ZIH/Traces/Vampir-Tutorial-Analysis-Exampler Process 26<	Process	14	sync		â				mm							Property		Value					
Process 16 Sync O <	Process	15	sync		<u>nn</u>				mm							 File 		/Users/n	naweber/ZI	H/Traces/\	/ampir-Tut	orial-Analy	sis-Example
Process 17 syme	Process	16	sync	0 0												Creator		VampirT	race 5.4b1				
Process 19 sync	Process	17	sync	0 0	$\overline{\mathbf{m}}$	500	nnn		m			mm				Version		1.3.8 jell	yfish				
Process 19 sync	Process	18	sync	6	000) on	mma		\mathbf{m}		10.01					Number of	Processes	64					
Process 20 sync C CODDOCODDOCODDOCODDOCODDOCODDOCODDOC	Process	19	sync	0 0	000	503	nnn		m		mm	mm	$\overline{\mathbf{n}}$	mm		Timer Res		382 701	875 ne				
Process 21 sync CODDOCODDOCODDOCODDOCODDOCODDOCODDOCOD	Process	20	sync	0	$\tilde{\mathbf{m}}$	non	nnn		min	mmm	mii			mmm		Timer Res	Jution	502.701	070 p3				
Process 22 sync C C C C C C C C C C C C C C C C C C C	Process	21	sync	0	000		nnn		min														
Process 23 sync CODDOCODDOCODDOCODDOCODOCODDOCODOCODOCO	Process	22	sync	0	000				m														
Process 24 sync CODOCODOCODOCODOCODOCODOCODOCODOCODOCOD	Process	23	sync	0 0	0001																		
Process 25 sync Control Contro	Process	24	sync	0 0	0001														Functio	● n Legend			
Process 26 sync Control Contro	Process	25	sync	0 0	om	<u>iioi</u> i										Apr	lication						
Process 27 sync CODDOCODDOCODDOCODDOCODDOCODDOCODDOCOD	Process	26	sync	0	000	ii oii				<u>ticici i i</u>						DY	N						
Process 28 sync CODDOCODDOCODDOCODDOCODDOCODDOCODDOCOD	Process	27	sync	•	000								000000	mom		IO_	NETCDF						
Process 29 sync CODDD-DDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDD	Process	28	sync	0	000)))))							MP							
Process 30 sync CODDDD CODDDD CODDDD CODDDD CODDD CODD CODD CODDD CODD CODD CODDD CODD	Process	29	sync	0	000	iioii))))					oppositi		PH'	rs						
WRF	Process	30	sync	0	000	DO			mi							VT_	API						
																WR	F						
					_	_																	

- Weather forecast code
 WRF
- Run with 64 processes
- Dynamical core: e.g., density, temperature, pressure, and winds in the atmosphere (DYN group)
- Physical parameterization:
 e.g., clouds, rain, and
 radiation (PHYS group)

03_wrf_deimos



- Problem: 25% MPI run time fraction during the iterations of the model
- Behaviour is constant throughout all iterations
- Question: Which user function causes the problem? And why?

03_wrf_deimos



03_wrf_deimos



Floating point exceptions

by floating point (FP)
exceptions in WRF
Counter
FR_FPU_EXCEPTIONS_
SSE_MICROTRAPS shows
FP exceptions

Load imbalance is caused

VI-HPS

03_wrf_deimos

MODULE_ADVECT_EM::ADVECT_SCALAR occurrences shown in yellow



- Zoom into one iteration
- Function invocations of MODULE_ADVECT_EM::ADVE
 CT_SCALAR (shown in yellow) match high floating point exception occurrences indicated by the counter at the bottom

04_sbmfd4_jugene

	H			Timeline	• 🖷	- 78 - 78			133.201 s	Alian Proc. Alian.
	0,s	20 s	40 s	60 s	80 s	100 s	120 s	All Processes, Accumu	lated Exclusive Time per Fun	ction Group
				•				400,000	s 200,000 s	0 s
Process 0	00	000000						615,197.076 s		MPI
Process 1	00) 0000				288,323.222 s	MP
Process 2	00				> 0000		0000 00		95,751.102 s	MP_UTIL
Process 3	00	<u></u>) 0000		0000 00		77,616.469 s	FD4
Process 4	00		ve Obe				0000 00		9,076.292 s	METEO
Process 5	00) 0000		0000 00		3,332.737 s	COUPLE
Process 6	00	00000	•) 0000		6666			
Process 7	00	00000					(1000) ee			
Process 8	00		• •)		0006 00			
Process 9	00		•) 000 0		0000 00			
Process 10	00	000000) (000 0		0000 00		•	
Process 11	00)) () ()		6666		nde pp		Context View	
Process 12	00		•		0000		0000		× i Trace	× +
Process 13	00	000000			0000		- 0000 PP	Property	Value	
Process 14	00	000000	* 💽		pood			File	/Users/maweber/7IH/T	races/Vampir-Tutorial-Analysis-Fa
Process 15	00	0000000	•		6666		0000 00	Creator	VomnirTraco E 9	
Process 16	00	00000	•				dood eb	Creator	Valiipii Trace 5.6	
Process 17	00	00000	•		> 0000	ee Oppo	diade po	Version	1.6.4 stingray	
Process 18	00	•Oppodo	•) oppe	📫 💽 🏚	- 6666 • •	Number of Processes	8,192	
Process 19	00	• • ••••••	•		o opoc		oppor ep	Timer Resolution	1.176471 ns	
Process 20	00	00000	•) obo c		dicide pp	Trace Times		
Process 21	00	00000	•) ()			Start	Fri Jan 22 21:15:35 20	10 (1264191335911609)
Process 22	00		•) dooc			Stop	Fri Jan 22 21:17:50 20	10 (126/191/706/8599)
Process 23	00		•) 0000			Stop		(12041314/0040333)
Process 24	00	• •• •••••	•		oppe	📫 🌒 🏟	0000 00		Eunction Legend	
Process 25	00		• •		o opoc				Function Legend	
Process 26	00	00000	• 🐽		>		0000 00			
Process 27	00		•		0000		0000 00	ED4		
Process 28	00		•		opoc	•• 🗘	0000 00	METEO		
Process 29	00		• 🐽	🛛 🔿 ooddada) 000 0	•• 💽💬	0000 00			
Process 30	00		•• 🐽) 600 0	ee 💽 ee	0000 00			

- Weather forecast code
 COSMO-SPECS
- Run with 8192 processes
- COSMO: weather model (METEO group)
- SPECS: microphysics for accurate cloud calculation (MP and MP_UTIL group)
- Coupling of both models done in COUPLE group
- Dynamic load balancing (FD4 group)

04_sbmfd4_jugene

1	1	1 illi	1	6 📑 🔢	1		1	• 💡 I		23 s	5 110,594 s
					Timeline					××	Function Summary
			40 s	60 s	80 క	6	100 s	120 s			All Processes, Accumulated Exclusive Time per Function Group
Process 15	-			C CHURLEN					bb		400,000 s 200,000 s 0 s
Process 25	9 -	- MARINA				JULLE			100	i i i i i i i i i i i i i i i i i i i	551,665.37 s (60.89%) MPI
Process 56 Process 52	o 5 -	-popp		0000000000000				0000000	00	Í	204,380.539 s (22.56%) MP
Process 75	8 -	-8999						000000000			95,751.102 s (10.57%) MP_UTIL
Process 10 Process 13	59 - 48 -	50666		• • • • • • • • • • • • • • • • • • • •		0000000		000000000	00		49,690.456 s (5.48%) FD4
Process 15	76 -	\mathbb{Q}_{2}							00		3,332.737 s (0.37%) COUPLE
Process 190	53 - 14 -										1,165.497 s (0.13%) METEO
Process 23	73 -							Ofference			
Process 25	78 -	计试试机						000000000			
Process 31	40 59 -	53333		00000000				- Celeseeeo	0000		
Process 35	17 -	23999	9900 0 90					occoccocco	00		
Process 37 Process 40	22 - 11 -	1888		o cobacco				000000000			
Process 42	78 -	Dodqa	ගණ ගණාම	o <u>o o o o o o o o</u> o		ာက္ကိုက္တက္ဆင္		 ••••••••••••••••••••••••••••••••••••	00		
Process 45 Process 47	73 - 47 -	Dodob	စစ္ေ၀၀၀ခုခ	00 086666666	ာ တစ္စစ္စာစု	0000000) Obe e	000600000	00		All Processes Accumulated Evolutive Time per Function
Process 49	29 -	Doddo	deo costa			000000	ୁ ପ୍ରଥିବିତ	000000	00		
Process 52 Process 54	01 - 03 -	Qeaqq	¢¢o ⊙ ¢¢		္ 👤 စိုစိုခိုင်			00000	00		320 871 545 c (35 42%)
Process 57	32 -	- ii kikiki				arrana 🗸			00		204 220 520 c (22 56%)
Process 58	95 - 15 -	-66666				0000000		0000000	00		92 370 658 s (10.2%)
Process 64	81 -	Dooco		0 00000000				000000000	00		92,570.000 5 (10.2%) Spinitu4_auvection_pd_in
Process 67	26 -	Dogog	6000000				00000	0000000	00		55,500.007 S (9.77%) MPI_Alireduce
Process 69 Process 712	55 - 22 -	Daead	eeo cone	ရင္ ၀၀ဖဖစ္ေရ		0000000) 🛈 🏟	0000000	00		50,632.754 \$ (6.10%) MPI_Reduce
Process 73	17 -	Qeoo	¢¢oco∎si¢	စုံစ ပစ်စွဲစွဲစုံစုံစုံ		000000		00000000	00		50,029.53 \$ (5.59%) MPI_Aligather
Process 75 Process 78	56 - 35 -		69000030	D 00000000		000000		00000000	00		38,998.19 \$ (4.26%) FD4_COUPLECOUPLE_PUT
Process 80	95 -		abooceab	୦୦ ୦୯ର୍ଖ୍ୟର୍ବ୍ଦ	o Obleo			00000000	2000		2/,/15.638 \$ (3.06%) MPI_Finalize
		lucius T	: 1	tion One of	•						5,025.364 \$ (0.55%) MPI_Isend
All Process	es, Exc	siusive i	ime per Fund	tion Group							4,264.022 \$ (0.47%) FD4_COUPLECOUPLE_GET
	100 %										3,380.444 s (0.37%) sbmfd4_removeative_meteors
	80 %	ML									3,332./3/ s (U.3/%) met_set
	00 /0			MPI	MPL			MPI			2,440.7/3 s (0.27%) FD4_GHOSTCTCOMM_EXCH
	60 %								MPI		2,069./1/ s (0.23%) FD4_PART_SFSFC_WEIGHT
	40 %										1,462.982 s (0.16%) MPI_Recv
		MP									1,068.303 s (0.12%) FD4_BALANCECE_READJUST
	20 %							MP			1,018.507 s (0.11%) org_runge_kutta
	0 %	_									804.359 s (0.09%) MPI_Scatter
				:			:	:		-	428.811 s (0.05%) FD4_BALANCEK_MIGRATION
30.25 s											

- Problem: Large MPI runtime fraction (>60%) during iterations
- Especially in MPI_Waitall and MPI_Allreduce
- Behaviour is constant throughout all iterations

VI-HPS



Large runtime variation in
 MPI Allreduce

Sorted profile reveals processes with small

MPI_Waitall timeshare

- Reason: Load imbalance in FD4_COUPLE_PUT and FD4_COUPLE_GET
- Most processes need to
 wait at MPI_Allreduce
 and MPI_Waitall
 (asynchronously)



Summary

- Performance flaws can lead to significant runtime overheads
- Use resources efficiently
- Analyze your code
- Performance analysis tools are there to help you

VI-HPS



http://www.vampir.eu

vampirsupport@zih.tu-dresden.de