King Abdullah University of Science and Technology

KAUST

Code profiling on Shaheen using Xprof and MPI_Trace.

Nicholas Allsopp (KAUST)

Content

Xprof: Hardware (CPU) performance

MPI_Trace: Message-passing performance

Xprof

Subroutine + line level profiling.

Using the IBM compilers, compile and link: -g –pg -qfullpath

The -g option can be used along with optimization, but sometimes the actual parent or grandparent may be off a line or two.

Run as usual.

Generates output files called gmon.out.0….n-1

Xprof a.out gmon.out.0

Xprof: Main Display

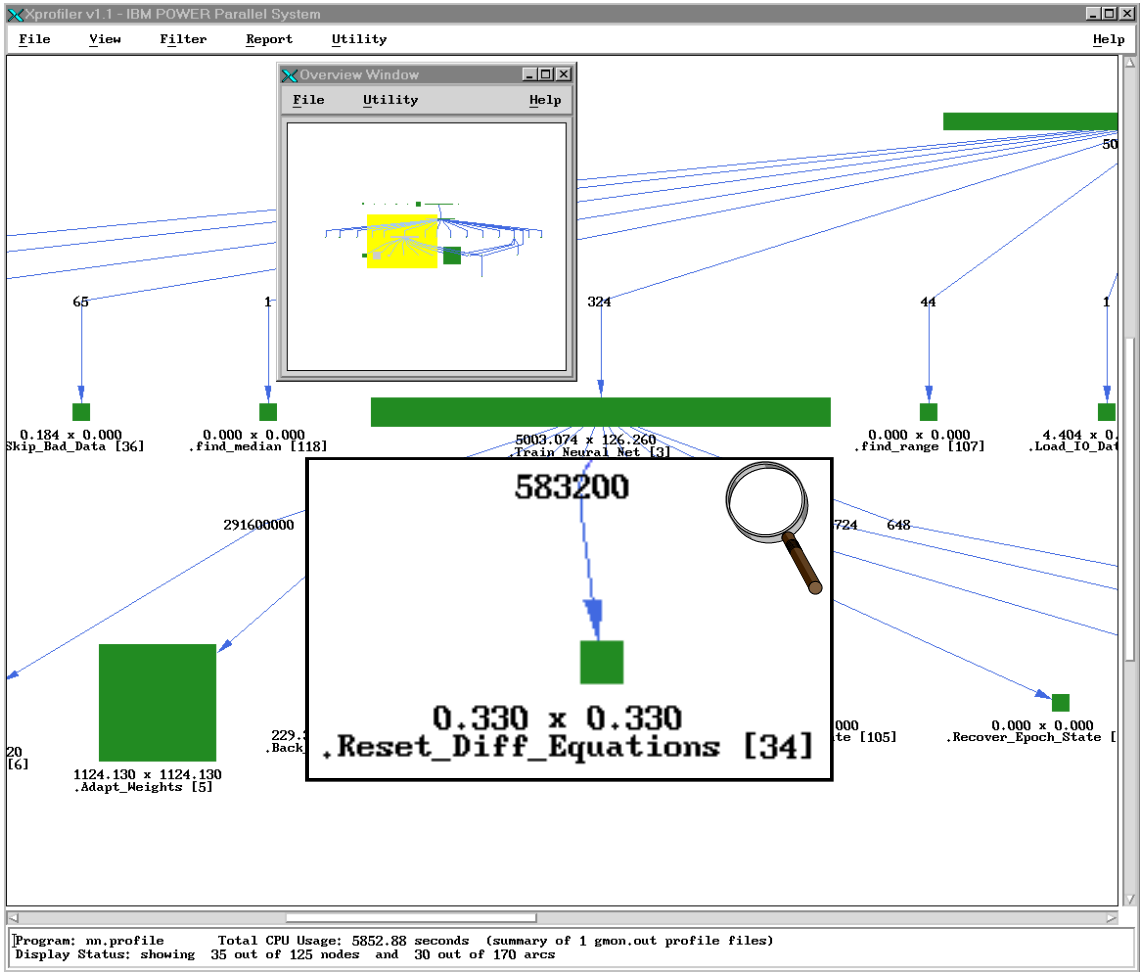

Width of a bar:
time including called routines

Height of a bar:
time excluding called routines

Call arrows labeled with number of calls

Overview window for easy navigation (View → Overview)

Xprof: Flat Profile

Menu **Report** provides usual gprof reports plus some extra ones

- Flat Profile
- Call Graph Profile
- Function Index
- Function Call Summary
- Library Statistics

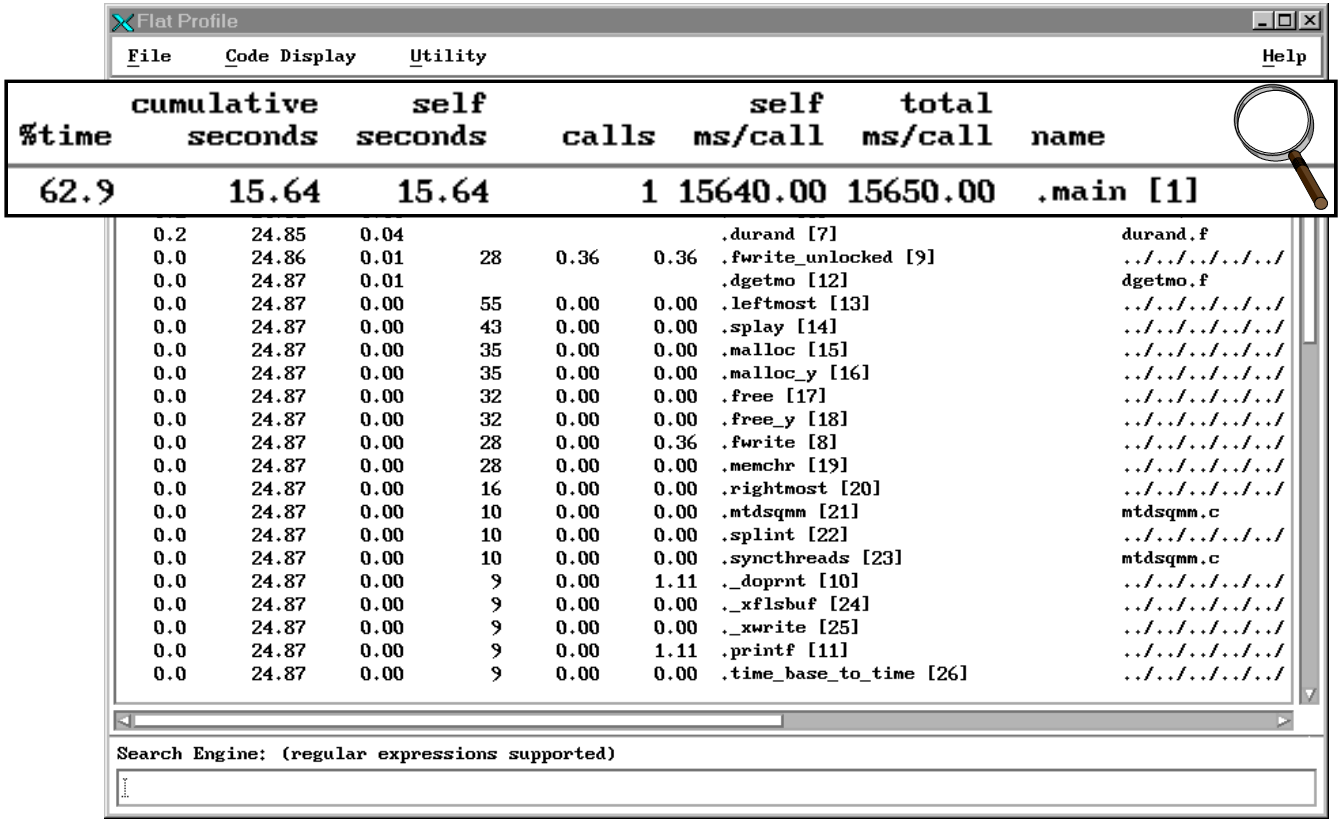

Xprof: Source Code Window

Source code
window displays
source code
with time profile
(in ticks=.01 sec)

Access

- Select function in main display
- → context menu
- Select function in flat profile
- → Code Display
- → Show Source Code

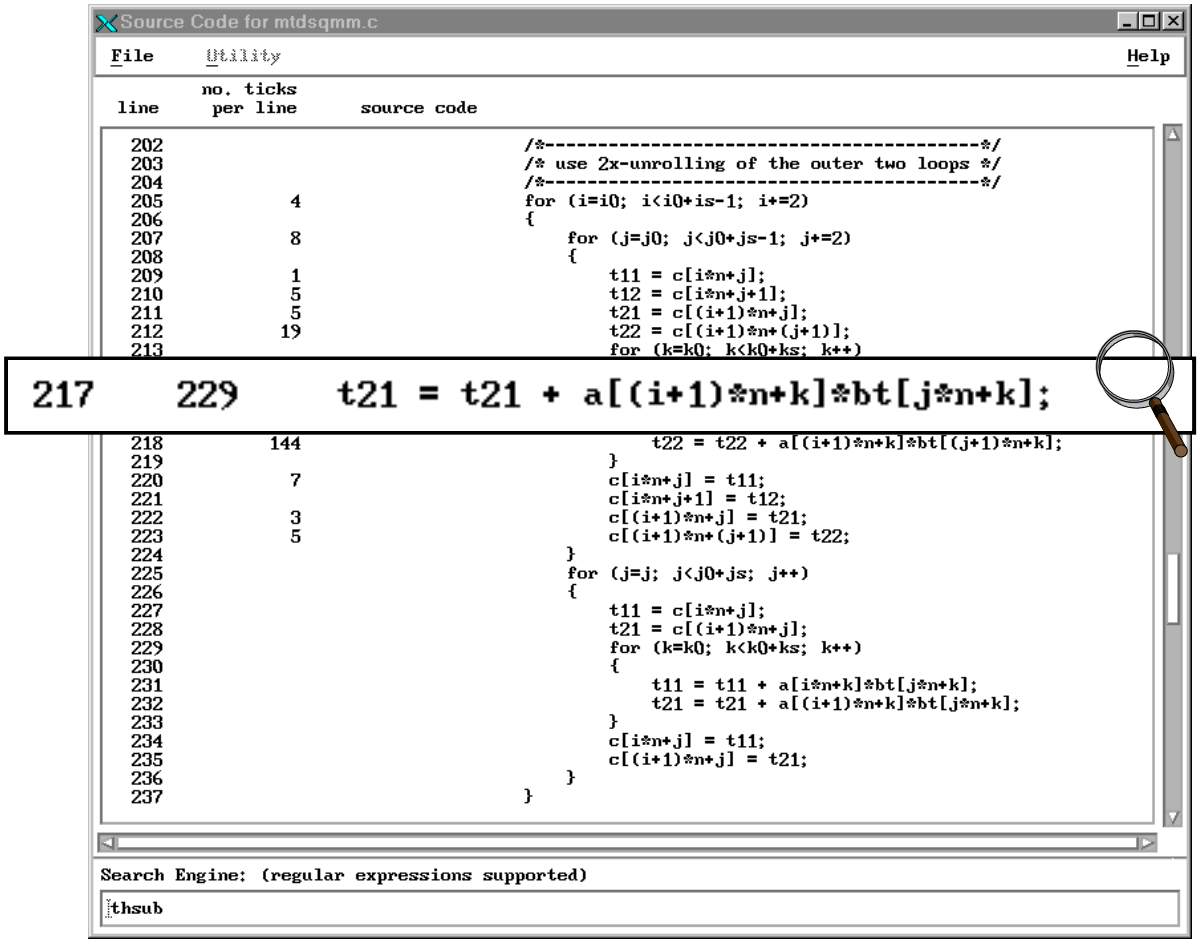

Xprof - Disassembler Code

Disassembler Code for .calc3 [3]

File Help

address	no. ticks per instr.	instruction	assembler code		source code
10002E18	81	FCC4287C	fnms	6,4,1,5	
10002E1C	64	CCF70008	lfdu	7,0x8(23)	POLD(I,J) = P(I,J)+ALPHA*(PNEW(I,J)-
10002E20	187	C90C0008	lfd	8,0x8(12)	
10002E24	53	C9750008	lfd	11,0x8(21)	UOLD(I,J) = U(I,J)+ALPHA*(UNEW(I,J)-
10002E28	89	FD63582A	fa	11,3,11	
10002E2C	63	FD28387C	fnms	9,8,1,7	POLD(I,J) = P(I,J)+ALPHA*(PNEW(I,J)-
10002E30	4	DD5B0008	stfdu	10,0x8(27)	U(I,J) = UNEW(I,J)
10002E34		C9540008	lfd	10,0x8(20)	VOLD(I,J) = V(I,J)+ALPHA*(VNEW(I,J)-
10002E38	113	FCCA302A	fa	6,10,6	
10002E3C	27	C8760008	lfd	3,0x8(22)	POLD(I,J) = P(I,J)+ALPHA*(PNEW(I,J)-
10002E40	87	FD8012FA	fma	12,0,11,2	UOLD(I,J) = U(I,J)+ALPHA*(UNEW(I,J)-
10002E44	35	DCB90008	stfdu	5,0x8(25)	V(I,J) = VNEW(I,J)
10002E48	4	FC63482A	fa	3,3,9	POLD(I,J) = P(I,J)+ALPHA*(PNEW(I,J)-
10002E4C	12	CD5A0008	lfdu	10,0x8(26)	UOLD(I,J) = U(I,J)+ALPHA*(UNEW(I,J)-
10002E50	62	FCC021BA	fma	6,0,6,4	VOLD(I,J) = V(I,J)+ALPHA*(VNEW(I,J)-
10002E54	36	C85B0008	lfd	2,0x8(27)	UOLD(I,J) = U(I,J)+ALPHA*(UNEW(I,J)-
10002E58	244	DCEC0008	stfdu	7,0x8(12)	P(I,J) = PNEW(I,J)
10002E5C	28	FD0040FA	fma	8,0,3,8	POLD(I,J) = P(I,J)+ALPHA*(PNEW(I,J)-
10002E60		C8990008	lfd	4,0x8(25)	VOLD(I,J) = V(I,J)+ALPHA*(VNEW(I,J)-
10002E64	316	DCD40008	stfdu	6,0x8(20)	
10002E68	29	FC62507C	fnms	3,2,1,10	UOLD(I,J) = U(I,J)+ALPHA*(UNEW(I,J)-

Search Engine: (regular expressions supported)

Message-Passing Performance: MP_Profiler Library

MP_Profiler

- Captures “summary” data for MPI calls with source code traceback
- No changes to source code, but MUST compile with -g
- ~1.7 microsecond overhead per MPI call
- Required link of mpitrace

Module load mpi_profile
Link in Environment variable: KSL_MPITRACE_LIB

Example of domain decomposition with single master I/O

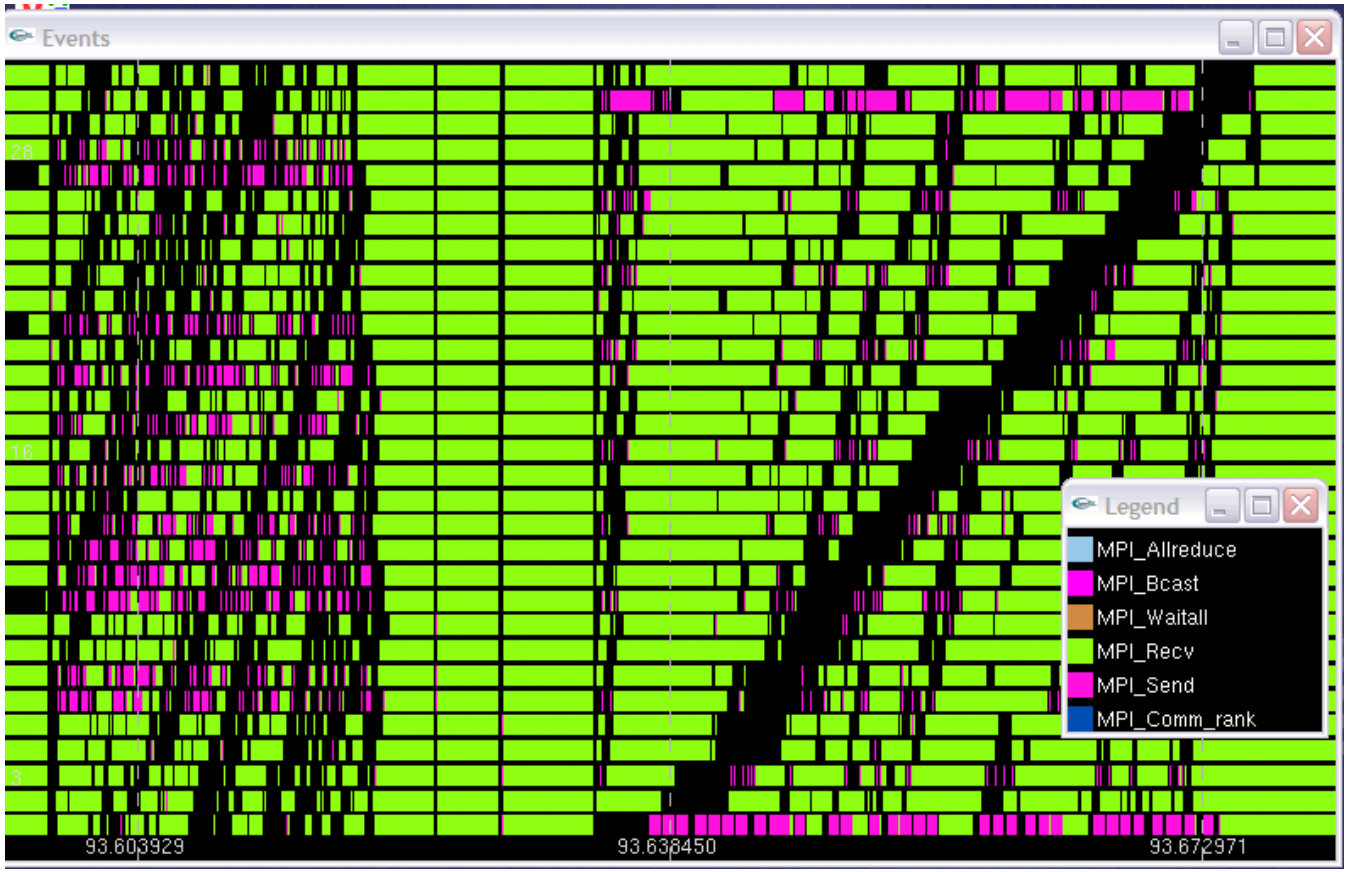

Swapping halo cells

Single node file output

Single node serialisation

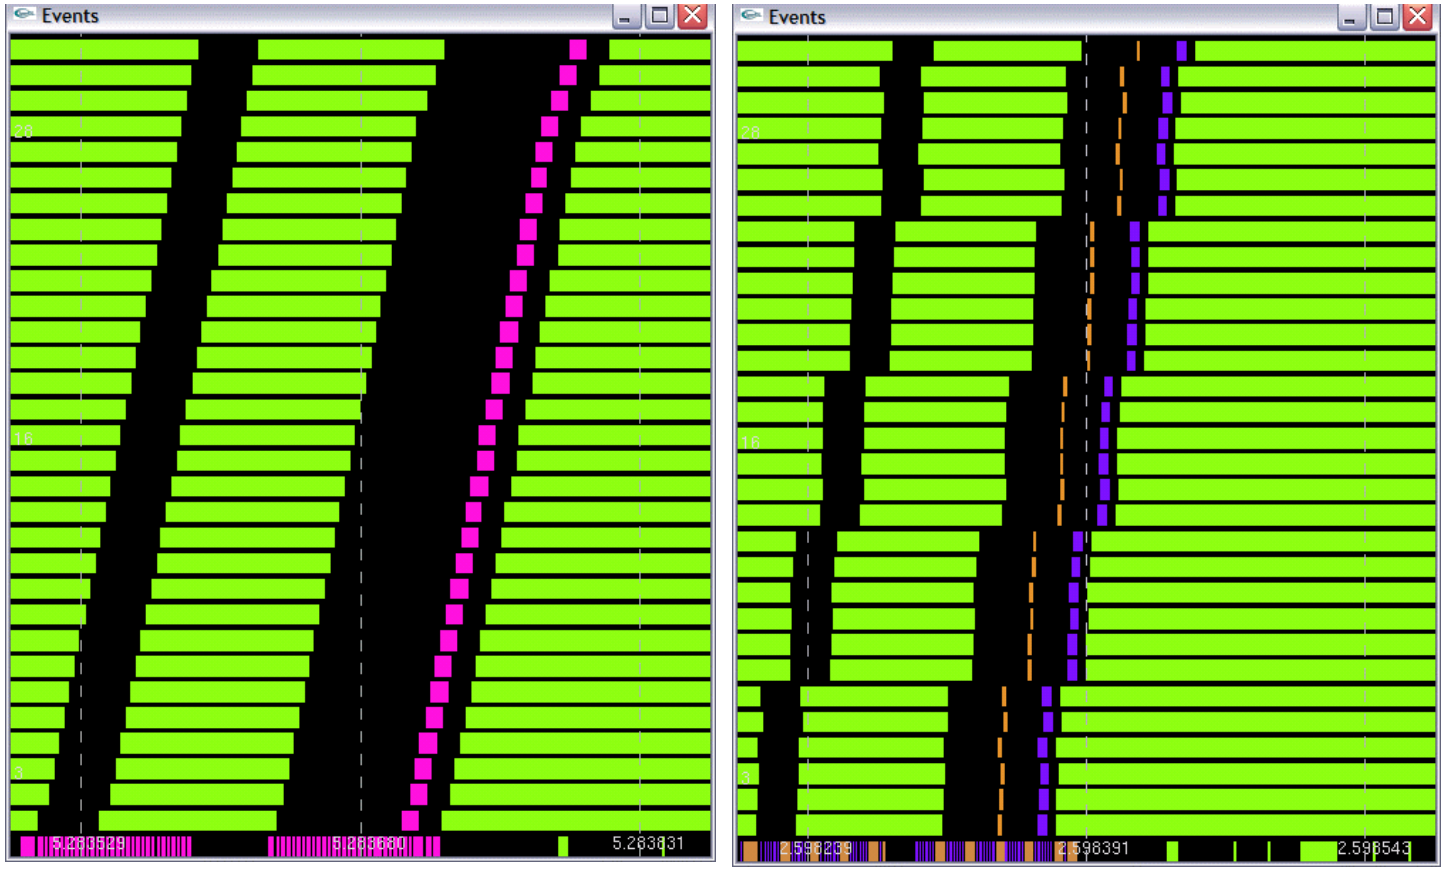

Using master to groups of workers. Time is along the x-axis and each row is a different MPI task. The master node is the bottom row

Remember when Using Xprof / MPI_TRACE

Using the IBM compilers, compile and link: -g –pg –qfullpath

Module load mpi_profile

Link in Environment variable: KSL_MPITRACE_LIB

Optional run-time environment variables:

-env "TRACE_ALL_EVENTS=yes TRACE_ALL_TASKS=yes SAVE_ALL_TASKS=yes”

To visualise:

Xprof a.out gmon.out.0

traceview.x events.trc

addr2line -e a.out hex_instruction_address

Questions