Sherpa W+Jet Profiling

Tim Martin, Warwick
June 19th 2020

Intro

- Noted in kickoff meeting that the largest CPU draw in ATLAS Event Generation comes from Sherpa SM processes. (CMS dominated by MadGraph, to a lesser extent)
- Investigate where this comes from, look for potential improvements.

➜ right plot: most CPU spent on high-precision calculations for $V + 0, 1, 2j$@NLO+$3, 4j$@LO and $t\bar{t} + 0, 1j$@NLO+$2, 3, 4j$@LO

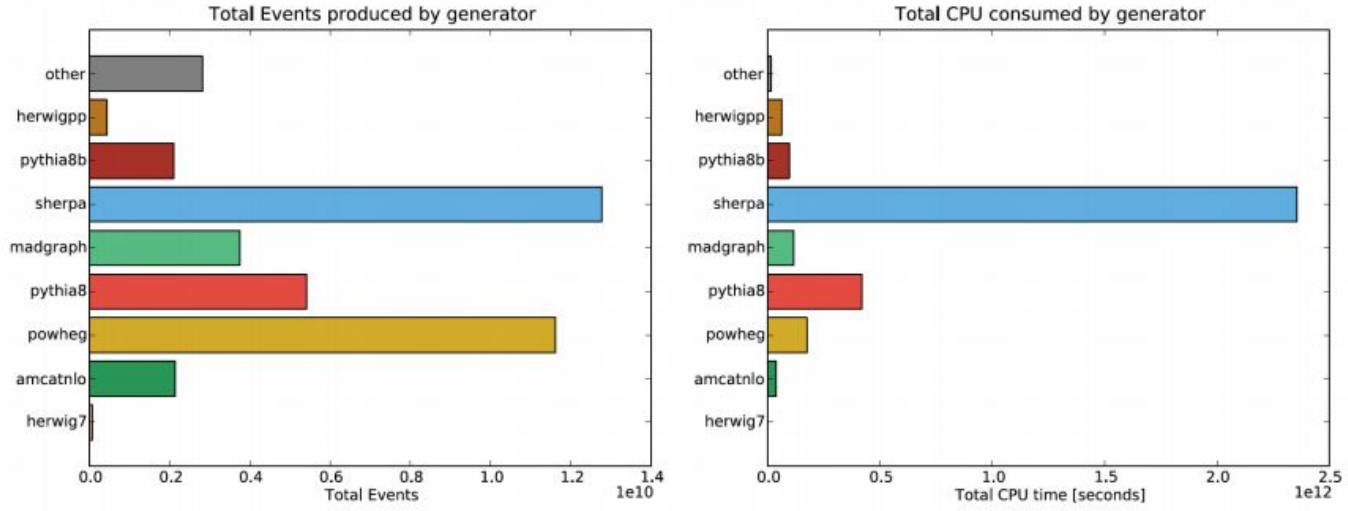

Source CHRISTIAN GÜTSCHOW ECHEP Feb Workshop

Initial State

- Everything compiled out-the-box at O2
- Full ATLAS-representative W+Jets setup provided by Marek Schoenherr
  - **W+0,1,2j@NLO+3,4,5j@LO**
  - Including approximate virtual corrections and reweightings to different PDFs and scales
- Running 500 events EvGen
- Single-core
- Total time: 19,876 s (around 5h 30m)

Software and PC Details

- Local compilations of
  - Sherpa 2.2.8
  - OpenLoops 2.1.1
  - LHAPDF 6.2.3
  - HepMC 3.2.0
  - Intel(R) VTune(TM) Profiler 2020 (build 605129)
- Local software stack provides
  - gcc 5.5.0
  - Intel(R) icc 16.0.3 20160415
  - cmake 3.12.2
  - sqlite 3.24.0
  - root 6.1

Profiling data collected by vtune running in userspace.
Visualisations of vtune database via **flamegraph**.

```
CPU Details
vendor_id : GenuineIntel
cpu family      : 6
model           : 158
model name      : Intel(R) Core(TM) i7-7700 CPU @ 3.60GHz
stepping  : 9
microcode : 0xd6
cpu MHz         : 900.000
cache size      : 8192 KB
```

Vanilla CPU #1: log

- Unsurprisingly large call on transcendental functions. Two main callees
  - `-> LHAPDF::LogBicubicInterpolator::_interpolateXQ2 -> LHAPDF::Interpolator::interpolateXQ2 -> LHAPDF::GridPDF::_xfxQ2 -> LHAPDF::PDF::xfxQ2 -> PDF::LHAPDF_CPP_Interface::GetXPDF 1857.171s 0ms 0usec libLHAPDF.so LHAPDF::LogBicubicInterpolator::_interpolateXQ2(LHAPDF::KnotArray1F const&, double, unsigned long, double, unsigned long) const [Unknown] 0x453b0`
  - `-> MODEL::One_Running_AlphaS::AlphaSLam -> MODEL::One_Running_AlphaS::operator() 1101.581s  0ms  0usec libModelMain.so.0  MODEL::One_Running_AlphaS::AlphaSLam(double, int)  Running_AlphaS.C  0x221b0`

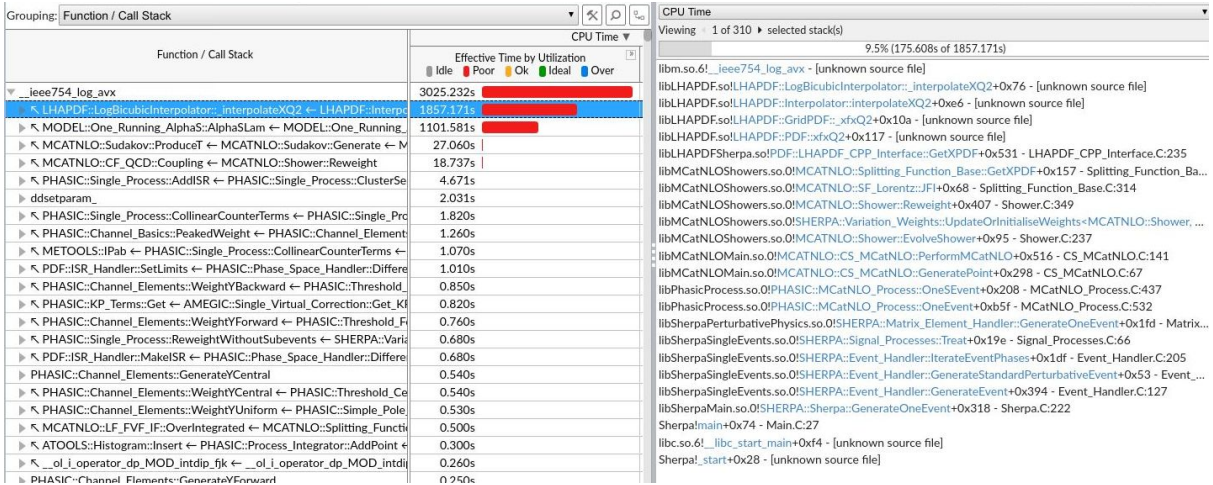

Calls to `__ieee754_log_avx`

Vanilla CPU #2: _interpolateXQ2

- As well as the time spent calling log, the function amasses a further 1,211 s

Vanilla CPU #3: exp

- Another transcendental takes the third slot at 1,174 s
- `MCATNLO::CF_QCD::CplFac  1174.721s  0ms  0usec libMCatNLOCalculators.so.0`
  `MCATNLO::CF_QCD::CplFac(double const&) const  CF_QCD.C  0xb420`

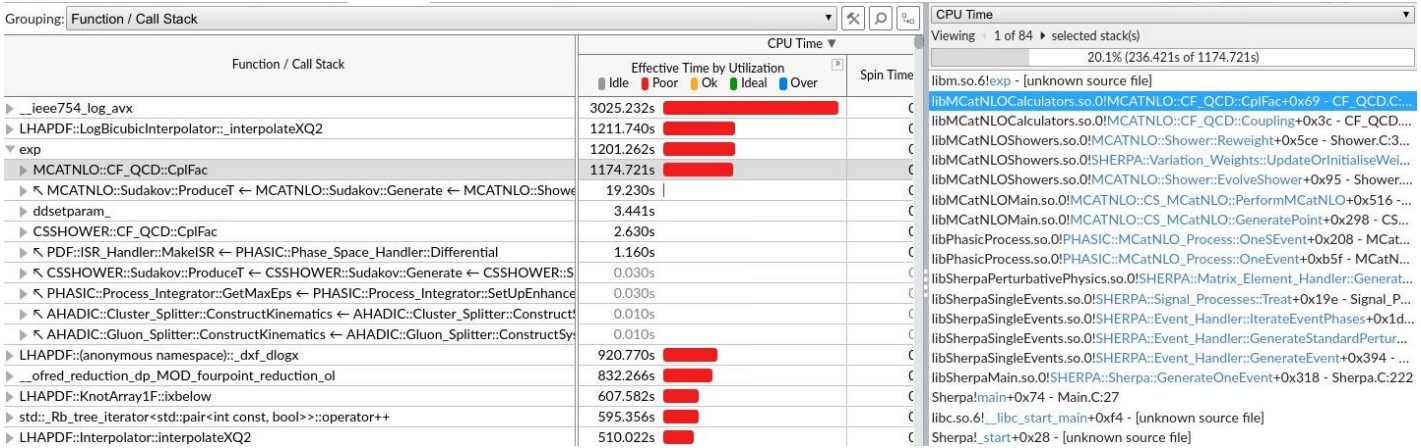


```
191  double CF_QCD::CplFac(const double &scale) const
192  {
193    if (m_kfmode==-1) return 1.0;
194    if (m_kfmode==0) return m_cplfac;
195    One_Running_AlphaS * const as = (p_altcpl) ? p_altcpl : p_cpl->GetAs();
196    double nf=as->Nf(scale);
197 >  double kfac=exp(-(67.0-3.0*sqr(M_PI)-10.0/3.0*nf)/(33.0-2.0*nf));   1.2%
198    return m_cplfac*kfac;
199  }
200
```


Flame Graph
Reset Search
Calls to __ieee754_exp_avx
~100% called by MCATNLO::CF_QCD::CplFac
libm..
libLHAPDF...
libLHAPDF.so!LH..
libLHAPDF.so!LHA..
libLHAPDF.so!LHAPD..
libLHAPDFSherpa.so!..
libMCatNLOShowers.so.0..
libMCatNLOShowers.so.0!MCATNLO::Shower::Reweight
libMCatNLOShowers.so.0!SHERPA::Variation_Weights::UpdateOrInitialiseWeights<..
libMCatNLOShowers.so.0!MCATNLO::Shower::EvolveShower
libMCatNLOMain.so.0!MCATNLO::CS_MCatNLO::PerformMCatNLO
libMCatNLOMain.so.0!MCATNLO::CS_MCatNLO::GeneratePoint
libPhasicProcess.so.0!PHASIC::MCatNLO_Process::OneSEvent
libPhasicProcess.so.0!PHASIC::MCatNLO_Process::OneEvent
libSherpaPerturbativePhysics.so.0!SHERPA::Matrix_Element_Handler::GenerateOneEvent
libSherpaSingleEvents.so.0!SHERPA::Signal_Processes::Treat
libSherpaSingleEvents.so.0!SHERPA::Event_Handler::IterateEventPhases
libSherpaSingleEvents.so.0!SHERPA::Event_Handler::GenerateStandardPerturbativeEvent
libSherpaSingleEvents.so.0!SHERPA::Event_Handler::GenerateEvent
libSherpaMain.so.0!SHERPA::Sherpa::GenerateOneEvent
Sherpa!main
libc.so.6!__libc_start_main
Sherpa!_start
[Unknown]!Total
all
libopenloop..
libopenloops_pplnj..
libopenloops.so!__..
libopenloops.so!__openloop..
libopenloops.so!ol_evaluate..
libSherpaOpenLoops.so!Ope..
libSherpaOpenLoops.so!OpenLoops:..
libDipoleSubtraction.so.0!AMEGIC::S..
libDipoleSubtraction.so.0!AMEGIC::Si..
libPhasicProcess.so.0!PHASIC::Single_Process::Differ..
libPhasicMain.so.0!PHASIC::Phase_Space_Handler::C..
libPhasicMain.so.0!PHASIC::Phase_Space_Handler::Differe..
libPhasicMain.so.0!PHASIC::Phase_Space_Handler::OneE..
libPhasicProcess.so.0!PHASIC::Process_Group::OneEvent
Matched: 5.6%

Vanilla CPU #4: _dxf_dlogx

- Like #1, #4 is a child of #2 **`_interpolateXQ2`**
- LHAPDF::LogBicubicInterpolator::_interpolateXQ2 -> LHAPDF::Interpolator::interpolateXQ2 -> LHAPDF::GridPDF::_xfxQ2 -> LHAPDF::PDF::xfxQ2 -> PDF::LHAPDF_CPP_Interface::GetXPDF 920.770s

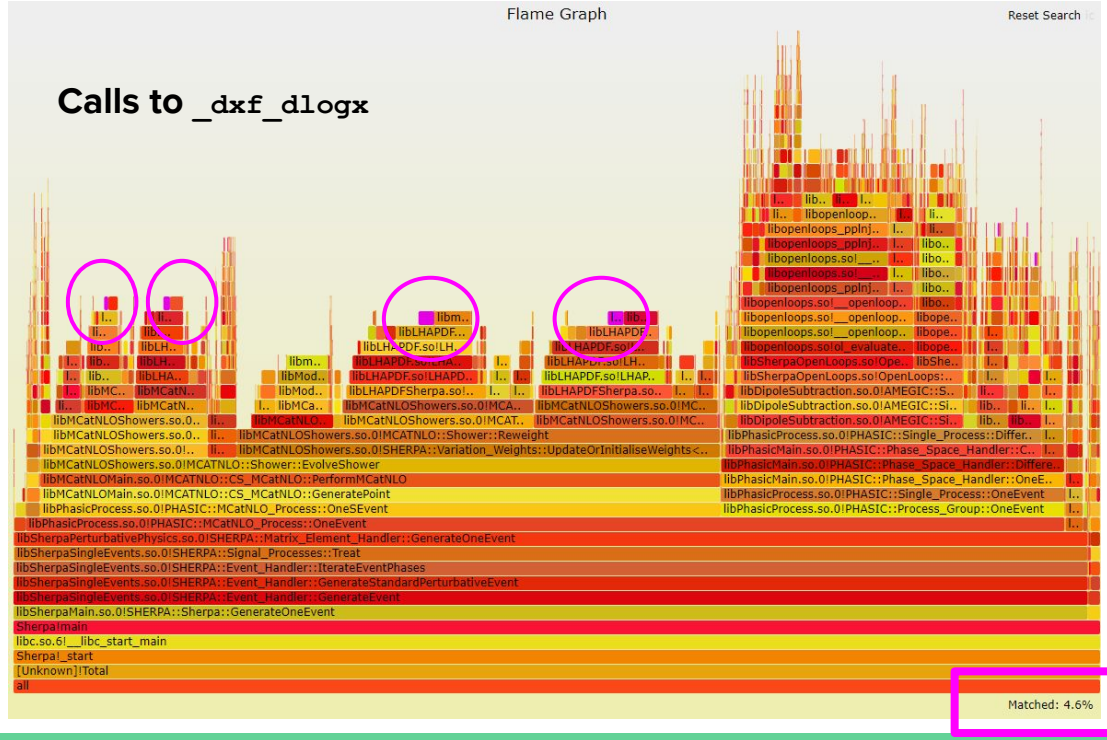

CPU Time

Viewing 1 of 1067 selected stack(s)

4.4% (40.866s of 920.770s)

libLHAPDF.so!LHAPDF::(anonymous namespace)::_dxf_dlogx - [unknown sour...
libLHAPDF.so!LHAPDF::LogBicubicInterpolator::_interpolateXQ2+0x120 - [un...
libLHAPDF.so!LHAPDF::Interpolator::interpolateXQ2+0xe6 - [unknown source...
libLHAPDF.so!LHAPDF::GridPDF::_xfxQ2+0x10a - [unknown source file]
libLHAPDF.so!LHAPDF::PDF::xfxQ2+0x117 - [unknown source file]
libLHAPDFSherpa.so!PDF::LHAPDF_CPP_Interface::GetXPDF+0x531 - LHAP...
libMCatNLOShowers.so.0!MCATNLO::Splitting_Function_Base::GetXPDF+0x1...
libMCatNLOShowers.so.0!MCATNLO::SF_Lorentz::JFI+0x68 - Splitting_Functi...
libMCatNLOShowers.so.0!MCATNLO::Shower::Reweight+0x407 - Shower.C:3...
libMCatNLOShowers.so.0!SHERPA::Variation_Weights::UpdateOrInitialiseWei...
libMCatNLOShowers.so.0!MCATNLO::Shower::EvolveShower+0x95 - Shower....
libMCatNLOMain.so.0!MCATNLO::CS_MCatNLO::PerformMCatNLO+0x516 -...
libMCatNLOMain.so.0!MCATNLO::CS_MCatNLO::GeneratePoint+0x298 - CS...
libPhasicProcess.so.0!PHASIC::MCatNLO_Process::OneSEvent+0x208 - MCat...
libPhasicProcess.so.0!PHASIC::MCatNLO_Process::OneEvent+0xb5f - MCatN...
libSherpaPerturbativePhysics.so.0!SHERPA::Matrix_Element_Handler::Generat...
libSherpaSingleEvents.so.0!SHERPA::Signal_Processes::Treat+0x19e - Signal_P...
libSherpaSingleEvents.so.0!SHERPA::Event_Handler::IterateEventPhases+0x1d...
libSherpaSingleEvents.so.0!SHERPA::Event_Handler::GenerateStandardPertur...
libSherpaSingleEvents.so.0!SHERPA::Event_Handler::GenerateEvent+0x394 - ...
libSherpaMain.so.0!SHERPA::Sherpa::GenerateOneEvent+0x318 - Sherpa.C:222
Sherpa!main+0x74 - Main.C:27
libc.so.6!__libc_start_main+0xf4 - [unknown source file]
Sherpa!_start+0x28 - [unknown source file]

Vanilla CPU #5: ofred_reduction_dp_MOD_fourpoint_reduction_ol

- Top single CPU calling coming from Openloops at 832 s

Vanilla O2 Summary

1. **98%** of time spent under **SHERPA::Matrix_Element_Handler::GenerateOneEvent**
   a. With **62.6%** spent under **PHASIC::MCatNLO_Process::OneSEvent**
      i. With **15.4%** spent under **MCATNLO::Shower::MakeKinematics**
      ii. With **44.5%** spent under **SHERPA::Variation_Weights::UpdateOrInitialiseWeights**
   b. With **31.6%** spent under **PHASIC::Process_Group::OneEvent**
      i. With **20.7%** spent under **AMEGIC::Single_Virtual_Correction::Partonic**

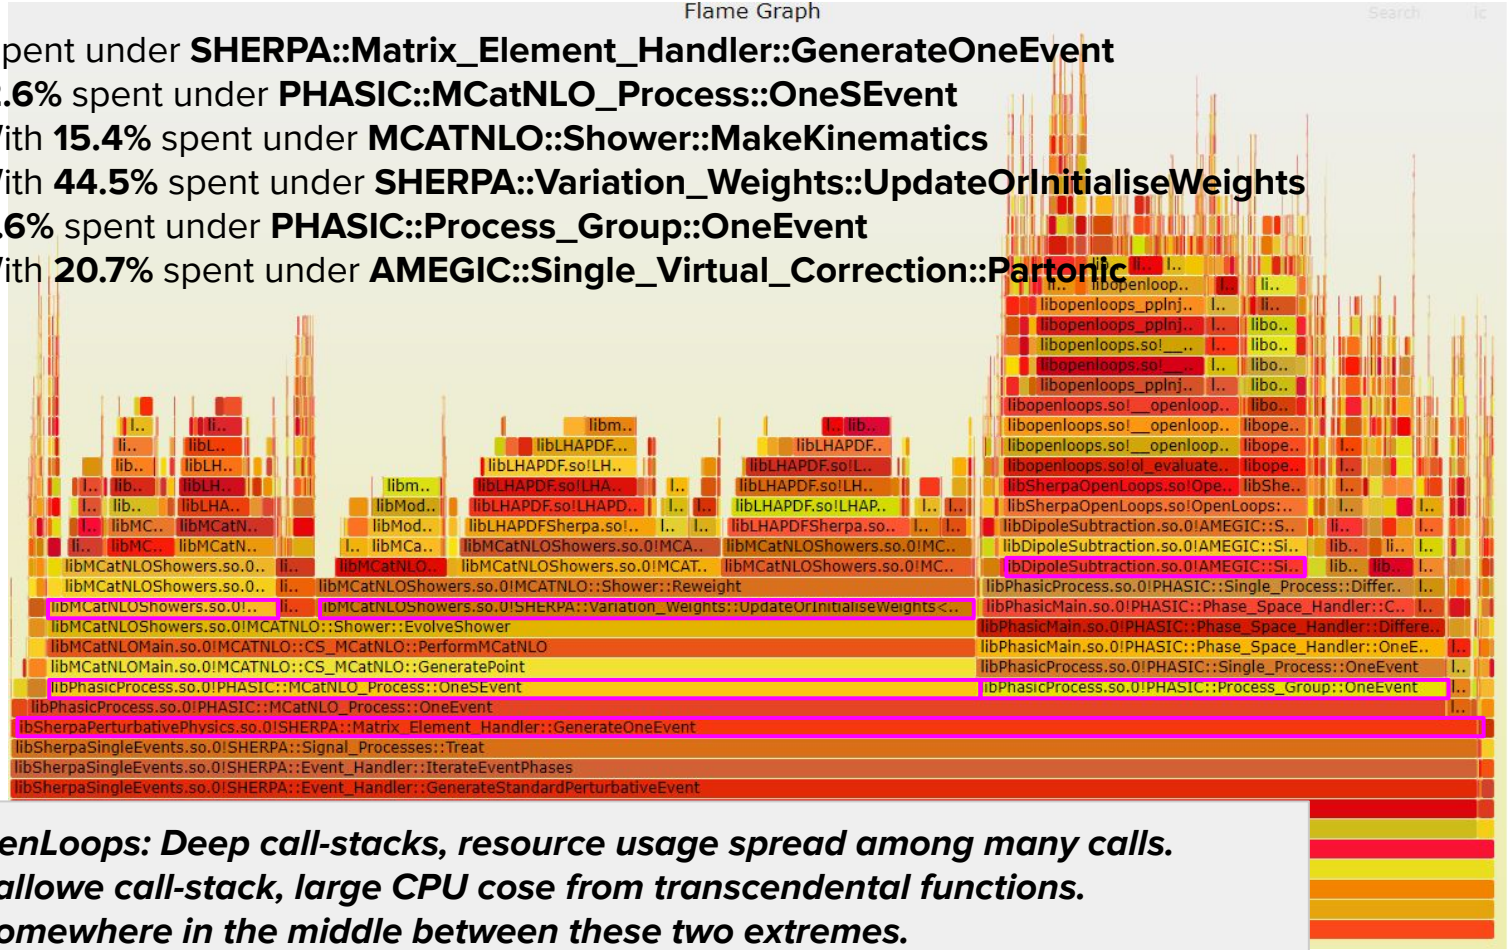

- ***AMEGIC, OpenLoops: Deep call-stacks, resource usage spread among many calls.***
- ***LHAPDF: Shallowe call-stack, large CPU cose from transcendental functions.***
- ***MCatNLO: Somewhere in the middle between these two extremes.***

Compile Time Optimisations

- Multiple easy **“slot in”** optimisation strategies tried.
  - **Memory allocation**: preload Google’s TCMALLOC
  - **Optimisation** level: O2 vs. O3
  - Beyond O3: **Minimum architecture** (msse4.2) and **unsafe maths** optimisations
  - **Link Time Optimisation**
  - Fully **static** single-process builds (**not managed successfully…**)
  - Use of Intel icc compiler.
    - O3, minimum architecture and unsafe maths flags with icc
- Optimisations were tried only on Sherpa and on Sherpa+LHAPDF. In some cases this resulted in a strange regression.
- TCMALLOC and VTune did not play nice.
- Working set of icc flags usable with both LHAPDF and Sherpa took some iterations.
- Have not checked so far if physics were impacted by unsafe maths.
- ***See backup slides for full details.***

Optimisation Timings

Optimisation	On Sherpa	On Sherpa+LHAPDF
None (O2)	19,876 s	N/A
Memory Allocation	19,505 s (-2%)	~83,463 s (!)
O3	19,506 s (-2%)	20,280 s (+2%)
Architecture & Unsafe Maths	20,183 s (+2%)	~82,785 s (!)
Link Time Optimisation	20,371 s (+2%)	~80,935 s (!)
Intel compiler		**18,122 s (-9%)**
Intel compiler + O3, Arch, Maths		18,598 s (-6%)

Intel savings look to come primarily through faster maths. C.f.

```
(gcc) __ieee754_log_avx = 3,025 s
(icc) __libm_log_l9 = 956 s
```

See also Manuel's slides from yesterday's ECHEP Reco meeting. Vectorisation can save much more.

LHAPDF Load

- May have noted, a lot of the areas noted so far are actually calls into the HEPMC library.
- Integrated, it sums to **47.3%** of all time.

Investigation into PDF Evaluation Calls

- Single event dumps of calls by Sherpa produced by Marek
  - <flavour>,<x>,<Q2> for some different types of events (pp->W+0j S/H, 1j S/H, 5j LO)
- Visualised by Andy

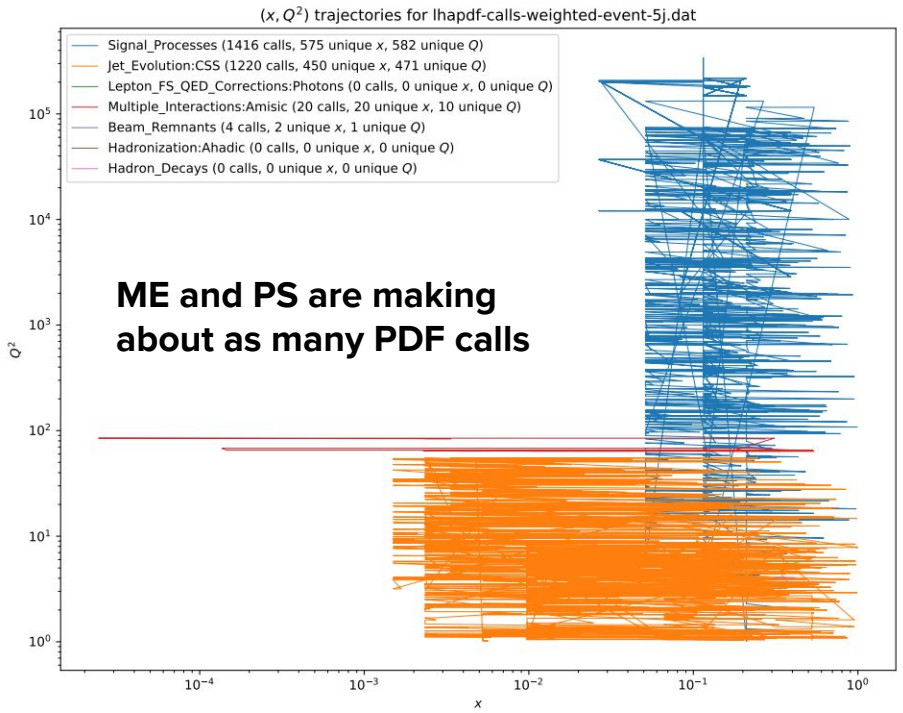

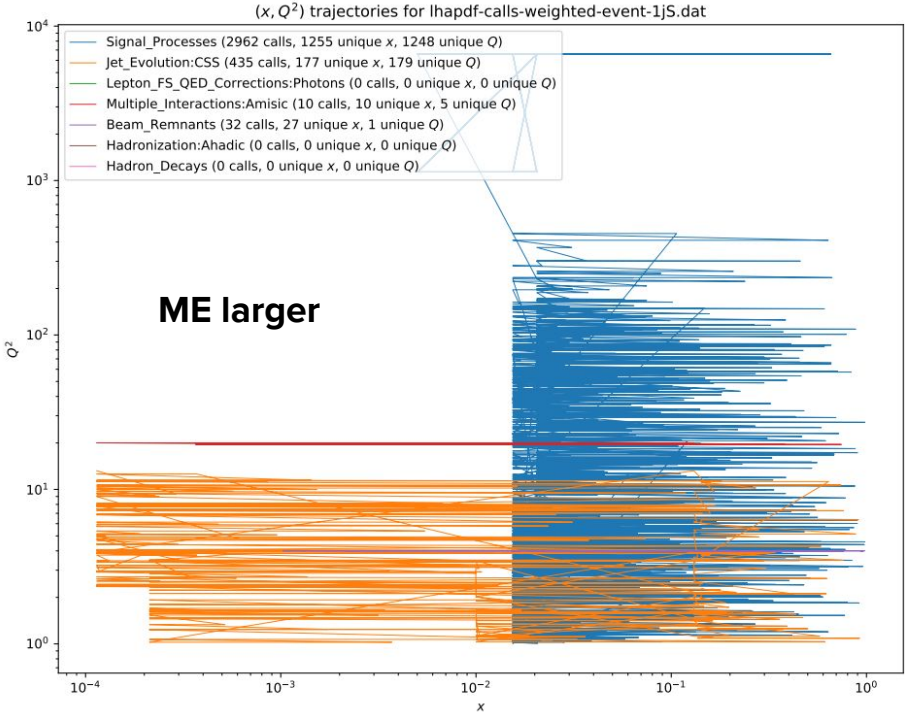

Caching in LHAPDF

Tested:
With 6.2.4beta1
With 6.3.0beta1
With 6.3.0beta3

- Investigating caching of values in LHAPDF
  - Some work carried out previously by Dima Konstantinov and Grigorii Latyshev
  - Additional work by Andy Buckley
- Iterations with implementing small thread-local caches **(current size: 4)**.
  - **Avoid re-interpolation** when the same value is requested multiple times in a single event.
- May also benefit from revisiting how Sherpa structures its LHAPFD calls.
- May also benefit from caching over different grids in LHAPDF.

Version	**% in LHAPDF**	**Total Time**
6.2.4	47.3%	19876 s
6.2.4beta1	43.9%	20235 s
6.3.0beta1	44.3%	20505 s
6.3.0beta2	43.7%	19634 s

6.3.0beta3

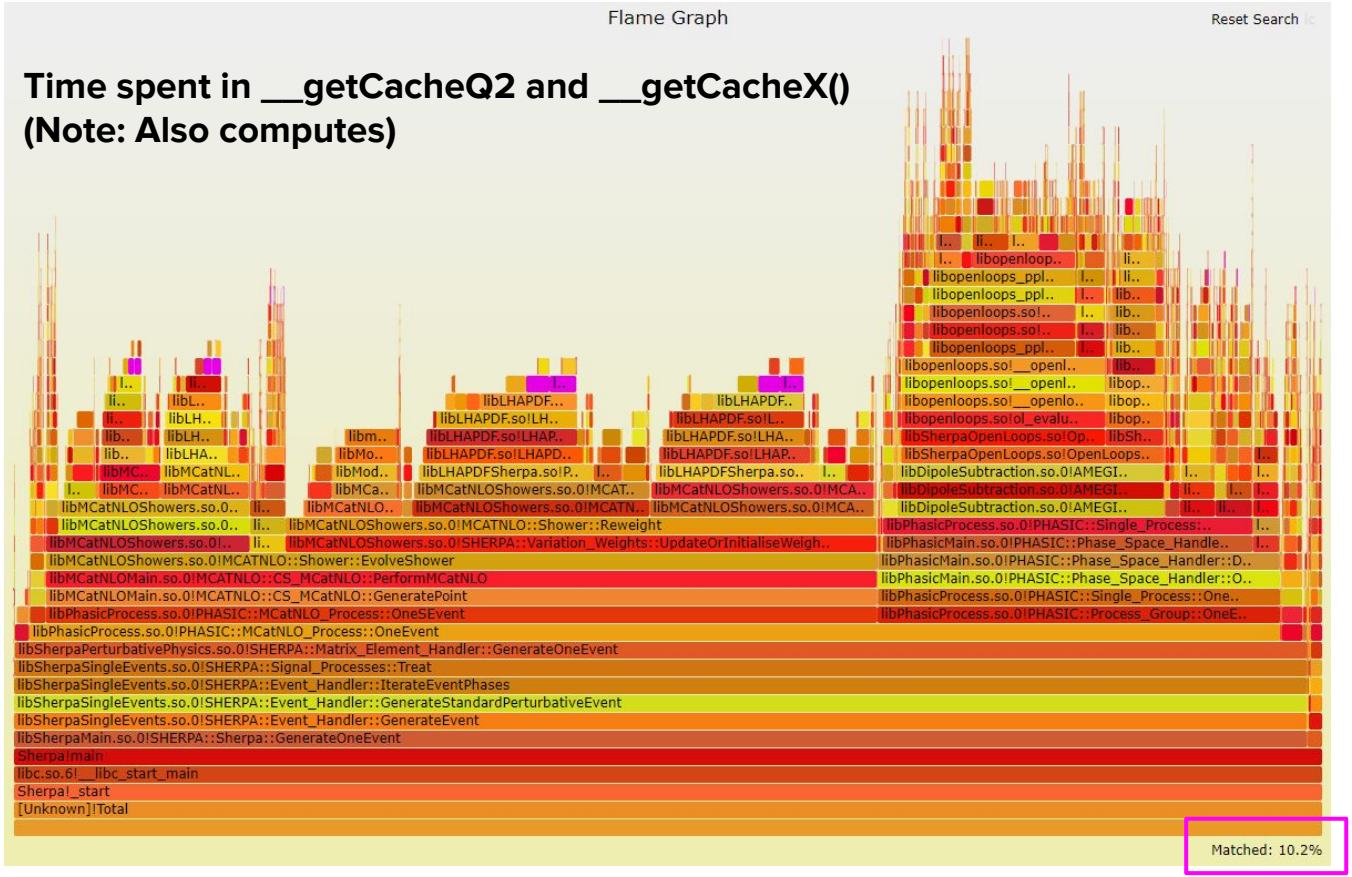
Flame Graph
Reset Search
Time spent in __getCacheQ2 and __getCacheX()
(Note: Also computes)
libLHAPDF..
libLHAPDF.so!LH..
libLHAPDF.so!L..
libLHAPDF.so!LHA..
libLHAPDF.so!LHAP..
libLHAPDFSherpa.so..
libMCatNLOShowers.so.0!MCAT..
libMCatNLOShowers.so.0!MCA..
libMCatNLOShowers.so.0!MCATN..
libMCatNLOShowers.so.0!MCATNLO::Shower::Reweight
libMCatNLOShowers.so.0!SHERPA::Variation_Weights::UpdateOrInitialiseWeigh..
libMCatNLOShowers.so.0!MCATNLO::Shower::EvolveShower
libMCatNLOMain.so.0!MCATNLO::CS_MCatNLO::PerformMCatNLO
libMCatNLOMain.so.0!MCATNLO::CS_MCatNLO::GeneratePoint
libPhasicProcess.so.0!PHASIC::MCatNLO_Process::OneSEvent
libPhasicProcess.so.0!PHASIC::MCatNLO_Process::OneEvent
libSherpaPerturbativePhysics.so.0!SHERPA::Matrix_Element_Handler::GenerateOneEvent
libSherpaSingleEvents.so.0!SHERPA::Signal_Processes::Treat
libSherpaSingleEvents.so.0!SHERPA::Event_Handler::IterateEventPhases
libSherpaSingleEvents.so.0!SHERPA::Event_Handler::GenerateStandardPerturbativeEvent
libSherpaSingleEvents.so.0!SHERPA::Event_Handler::GenerateEvent
libSherpaMain.so.0!SHERPA::Sherpa::GenerateOneEvent
SherpaMain
libc.so.6!__libc_start_main
Sherpa!_start
[Unknown]!Total
libopenloop..
libopenloops_ppl..
libopenloops.so!..
libopenloops_ppl..
libopenloops.so!..
libopenloops.so!..
libopenloops_ppl..
libopenloops.so!__openl..
libopenloops.so!__openl..
libopenloops.so!__openlo..
libopenloops.so!ol_evalu..
libSherpaOpenLoops.so!Op..
libSherpaOpenLoops.so!OpenLoops..
libDipoleSubtraction.so.0!AMEGI..
libDipoleSubtraction.so.0!AMEGI..
libDipoleSubtraction.so.0!AMEGI..
libPhasicProcess.so.0!PHASIC::Single_Process:..
libPhasicMain.so.0!PHASIC::Phase_Space_Handle..
libPhasicMain.so.0!PHASIC::Phase_Space_Handler::D..
libPhasicMain.so.0!PHASIC::Phase_Space_Handler::O..
libPhasicProcess.so.0!PHASIC::Single_Process::One..
libPhasicProcess.so.0!PHASIC::Process_Group::OneE..
Matched: 10.2%

Summary

- Sherpa W+jets is a known huge CPU consumer on ATLAS
- CPU savings can be made by use of Intel's maths library, more may be possible via vectorized maths libraries.
- Better integration with LHAPDF could yield larger saving still.
- Flamegraphs at https://cernbox.cern.ch/index.php/s/9OK2W17XELp6Gmt

Backup

Optimisations #1: malloc

- **`malloc`** consumes 1.7% and **`operator new()`** consumes 0.9%
- TCMalloc : Thread-Caching Malloc
- *“TCMalloc is Google's customized implementation of C's malloc() and C++'s operator new used for memory allocation within our C and C++ code. TCMalloc is a fast, multi-threaded malloc implementation.”*
- Around 2.3x faster than **`malloc`**, used by ATLAS
- Unfortunately, TCMalloc did not play nice with VTune. Just timed with **time**
- **Sherpa + TCMalloc: 19,505 s (-2%)**

Optimisations #2: -O3

- Effects of turning on higher levels of compiler optimisation in Sherpa.
- Very little change to call-graph, and to total time.
- **Sherpa -O3: 19,506s (-2%).**

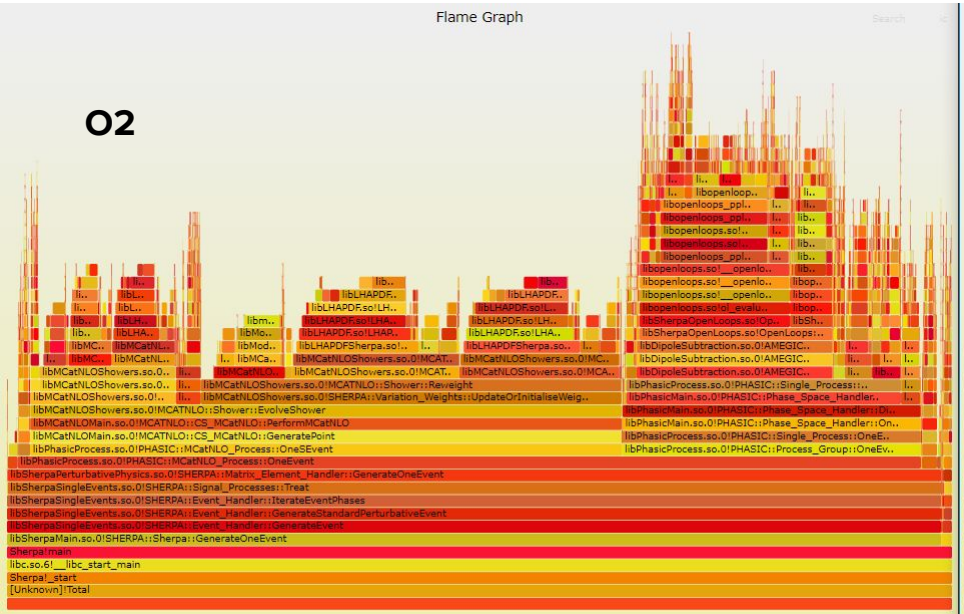

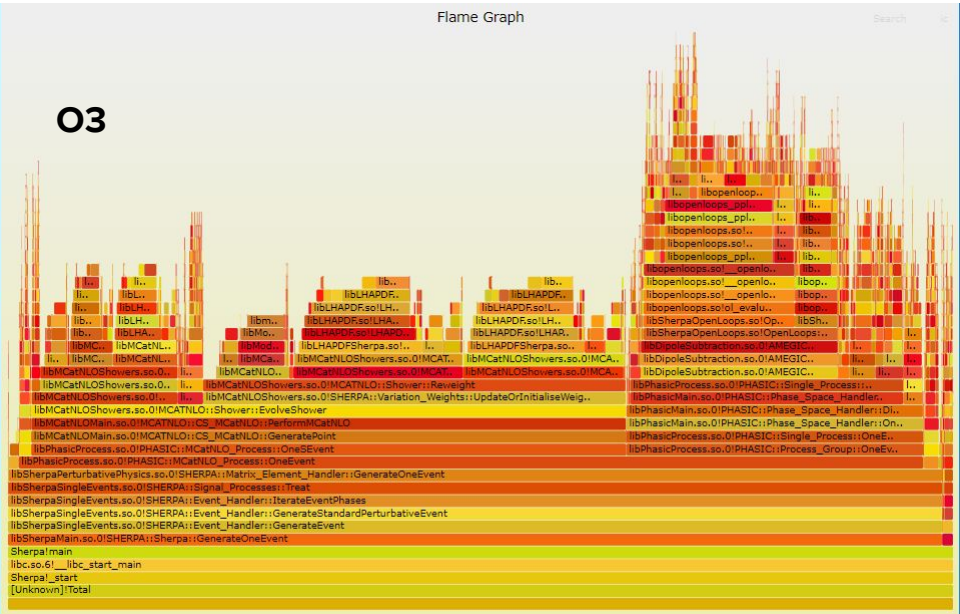

**IEEE 754 Violating*

Optimisations #3: Architecture & Unsafe Maths

- **`-msse4.2`** specifies to use the SSE4.2 CPU instruction set extension.
  - Available since Nehalem (initial core i5, i7), November 2008
  - LHCb reported having fully switched over to this instruction set.
  - ATLAS has recent experience of some grid nodes and user machines *still* not supporting it.
- **`-ffast-math`** applies a bunch of "unsafe"* operations, not applied by -O.
  - **`-fno-trapping-math, -fno-signaling-nans`**: *User **cannot trap /0** or overflow.*
  - **`-funsafe-math-optimizations`**: *This mode enables optimizations that allow arbitrary reassociations and transformations with no accuracy guarantees. Due to roundoff errors the associative law of algebra do not necessary hold for floating point numbers and thus **expressions like (x + y) + z are not necessary equal to x + (y + z)***
  - **`-ffinite-math-only`**: *Assume that there will **never be NaNs or +-Infs***
  - **`-fno-errno-math`**: *Disables setting of the errno variable as required by C89/C99 on calling math library routines.*
  - **`-fno-rounding-math`**: *IEEE has four rounding modes. This flag assumes that the rounding mode is round to nearest.*
  - **`-fcx-limited-range`**: *Causes the **range reduction step to be omitted when performing complex division**. This uses a / b = ((ar*br + ai*bi)/t) + i((ai*br - ar*bi)/t) with t = br*br + bi*bi and might not work well on arbitrary ranges of the inputs.*
  - **`-fno-signed-zeros`**: ***Removes the ability to have signed 0***

Optimisations #3: Architecture & Unsafe Maths

- Again, few changes to the call graph.
- Physics output not yet checked!
- **Sherpa -msse4.2 -ffast-math: 20,839 s (+5%)**

Needs re-running, other cores were in use…

Optimisations #4: Link Time Optimisation

- `-lto`: *Link time optimization is implemented as a GCC front end for a bytecode representation of GIMPLE that is emitted in special sections of .o files.*
- Additional data are generated by the compiler and passed via the .o files to the linker in order to allow it to perform additional optimisaiton.
- Total time: 20,371s (+2%)

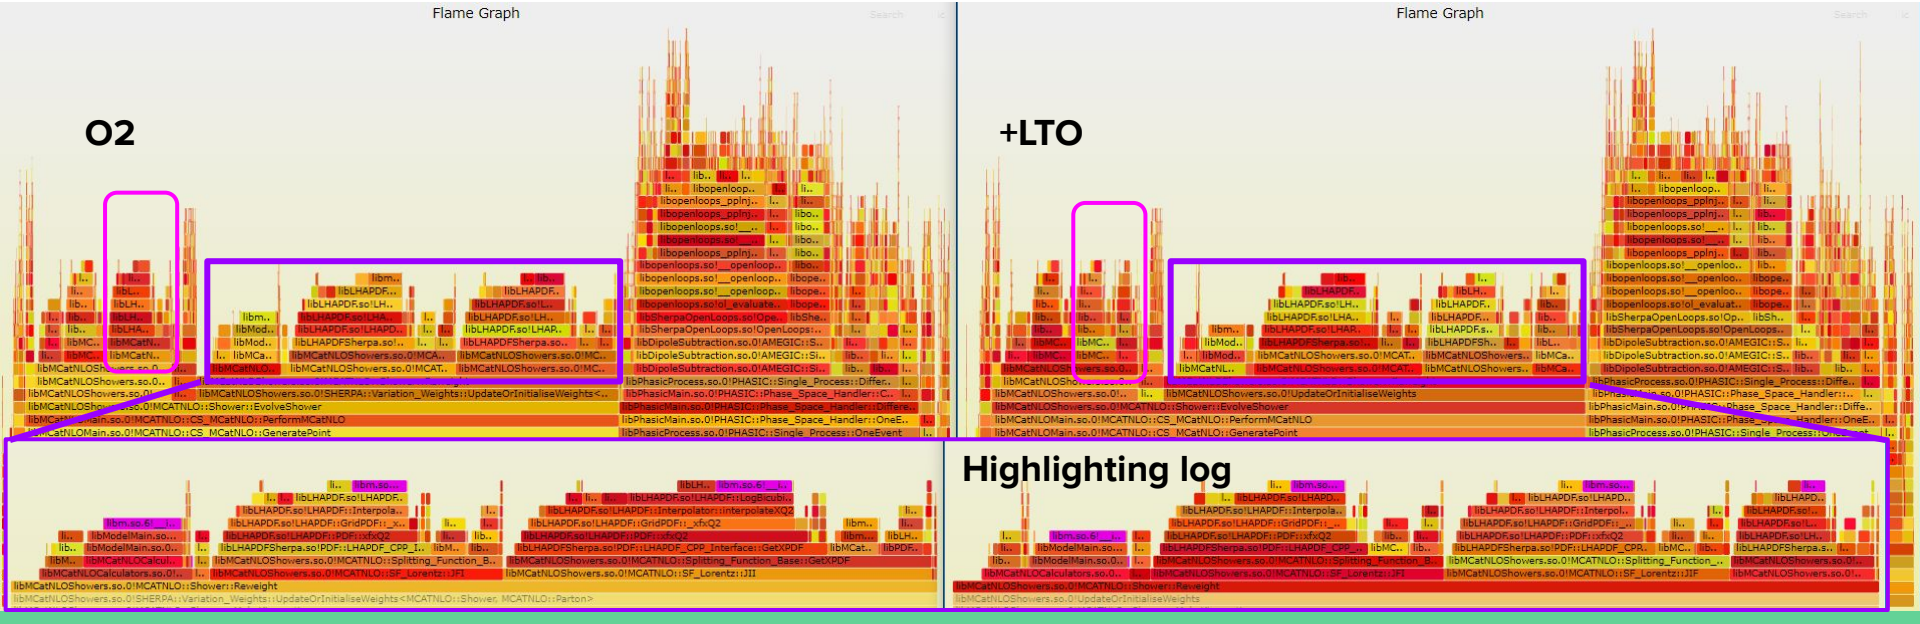

Other permutations:

- Trying to spread the flags also to LHAPDF encountered a serious regression!
  - O2: **40 s/event**
  - Sherpa O3 + LHAPDF O3 + OPENLOOPS fortran O3: **40 s/event**
  - Sherpa + TCMalloc , LHAPDF + TCMalloc: **167 s/event (!!!)**
  - Sherpa LTO + LHAPDF LTO: **161 s/event (!!!), 150 s/event (!!!)**
  - Sherpa Fast Maths flags + LHAPDF Fast Maths flags: **152 s/event (!!!)**

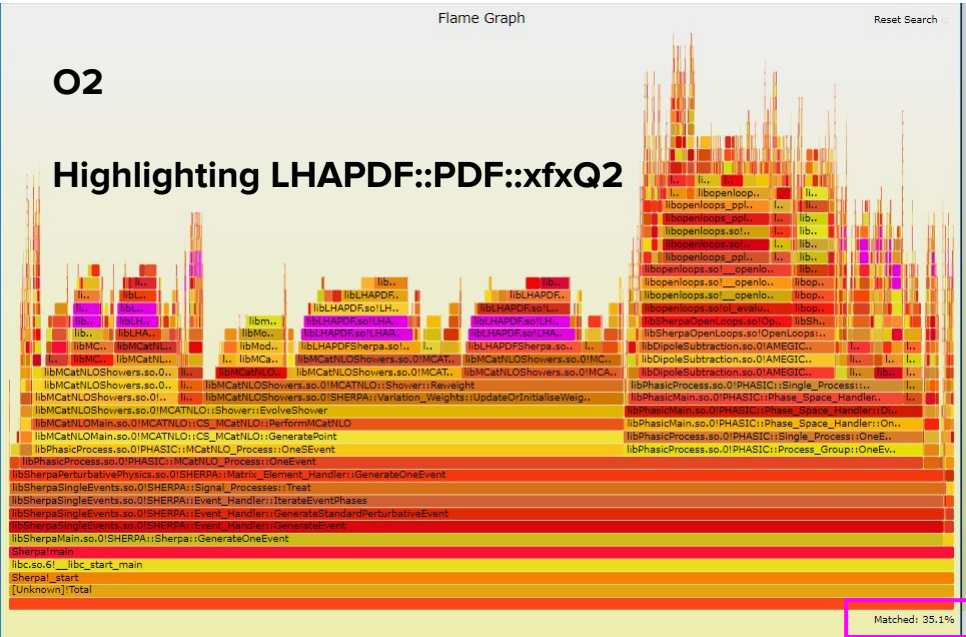

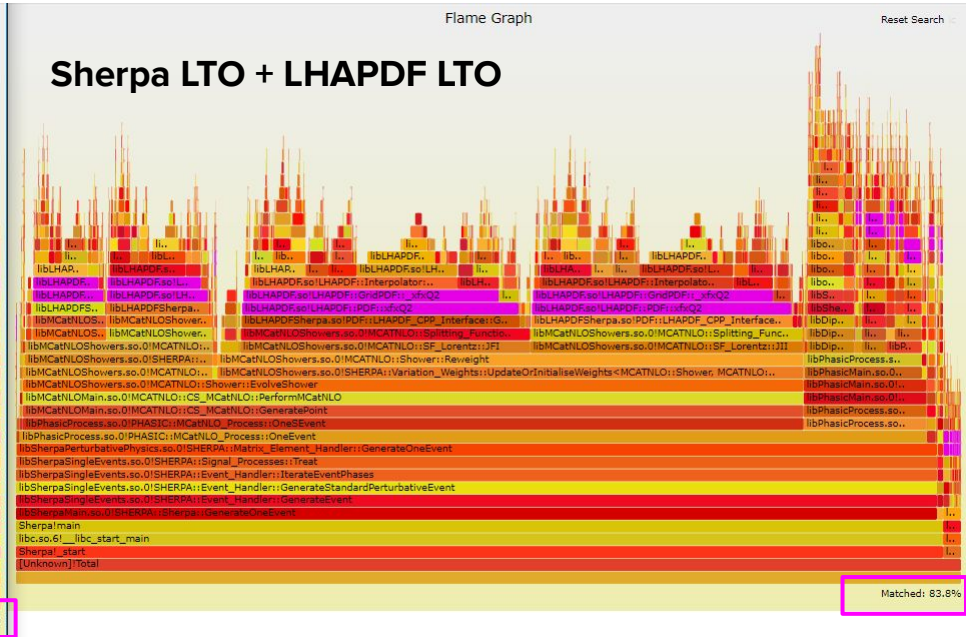

Other permutations:

- Trying a fully static build of Sherpa.
  - Tried to link Sherpa against minimal set of dependencies.
  - Build / combine static `.a` files for all required libraries.
  - Link them all together statically
- Managed to make a giant binary.
- Couldn't get it to work... segfaults...

Intel

- Proprietary compiler, but licencing may be available through CERN
- **Note**: ATLAS preload the intel maths libraries
- Using `icc 16.0.3 20160415` licenced by Warwick
- Total time: **18,025s**
  - **10% faster than GCC**

CPU #1: LHAPDF::LogBicubicInterpolator::_interpolateXQ2

- Predominantly split over `SF_Lorentz:JFI, JIF, JII`

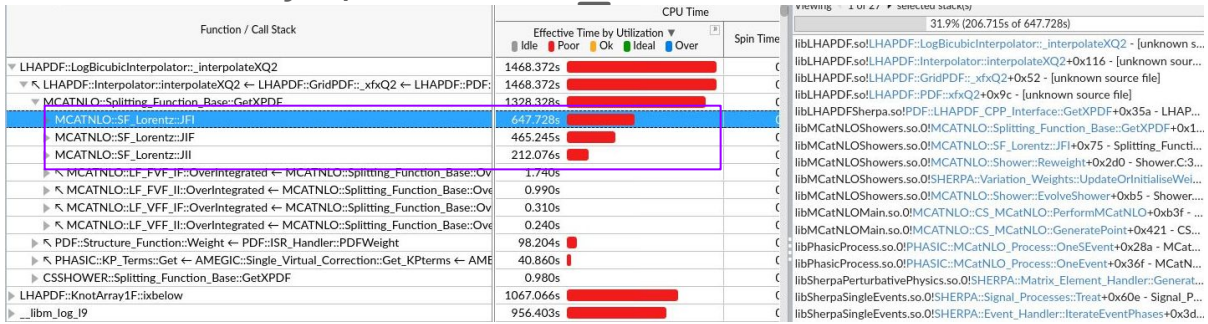

CPU #2: LHAPDF::KnotArray1F::ixbelow

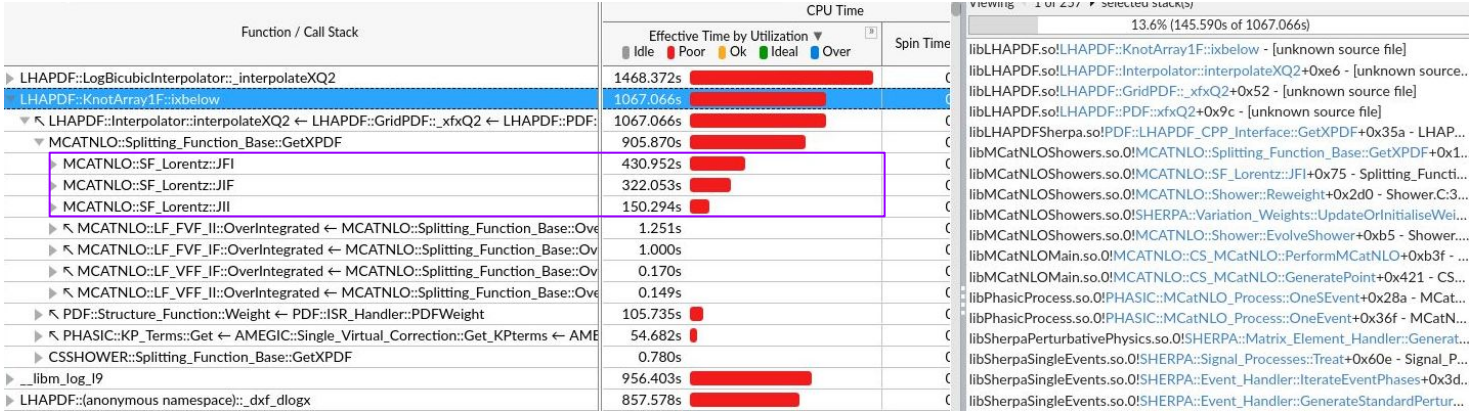

CPU #3: __libm_log_l9

- Log takes the #3 slot, calls from
  - LHAPDF::Interpolator::interpolateXQ2
  - MODEL::One_Running_AlphaS::operator()

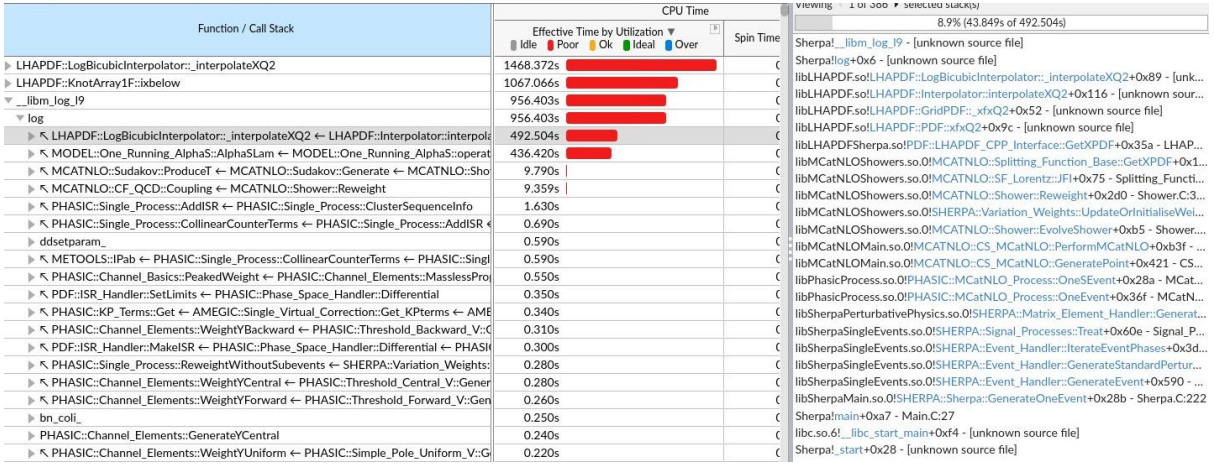

CPU #4: _dxf_dlogx

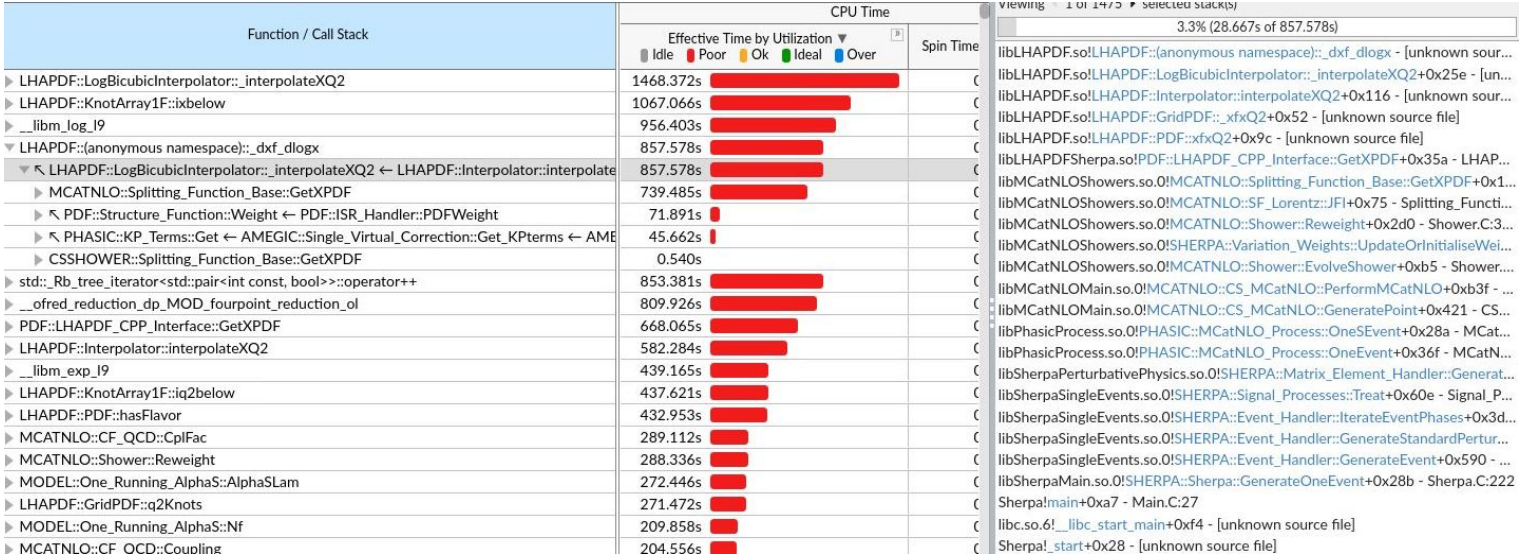

Function / Call Stack	CPU Time: Effective Time by Utilization (Idle / Poor / Ok / Ideal / Over)	CPU Time: Spin Time
LHAPDF::LogBicubicInterpolator::_interpolateXQ2	1468.372s	
LHAPDF::KnotArray1F::ixbelow	1067.066s	
__libm_log_l9	956.403s	
LHAPDF::(anonymous namespace)::_dxf_dlogx	857.578s	
↖ LHAPDF::LogBicubicInterpolator::_interpolateXQ2 ← LHAPDF::Interpolator::interpolate	857.578s	
MCATNLO::Splitting_Function_Base::GetXPDF	739.485s	
↖ PDF::Structure_Function::Weight ← PDF::ISR_Handler::PDFWeight	71.891s	
↖ PHASIC::KP_Terms::Get ← AMEGIC::Single_Virtual_Correction::Get_KPterms ← AME	45.662s	
CSSHOWER::Splitting_Function_Base::GetXPDF	0.540s	
std::_Rb_tree_iterator<std::pair<int const, bool>>::operator++	853.381s	
__ofred_reduction_dp_MOD_fourpoint_reduction_ol	809.926s	
PDF::LHAPDF_CPP_Interface::GetXPDF	668.065s	
LHAPDF::Interpolator::interpolateXQ2	582.284s	
__libm_exp_l9	439.165s	
LHAPDF::KnotArray1F::iq2below	437.621s	
LHAPDF::PDF::hasFlavor	432.953s	
MCATNLO::CF_QCD::CplFac	289.112s	
MCATNLO::Shower::Reweight	288.336s	
MODEL::One_Running_AlphaS::AlphaSLam	272.446s	
LHAPDF::GridPDF::q2Knots	271.472s	
MODEL::One_Running_AlphaS::Nf	209.858s	
MCATNLO::CF_QCD::Coupling	204.556s	

Viewing 1 of 1475 ▸ selected stack(s)

3.3% (28.667s of 857.578s)

libLHAPDF.so!LHAPDF::(anonymous namespace)::_dxf_dlogx - [unknown sour...
libLHAPDF.so!LHAPDF::LogBicubicInterpolator::_interpolateXQ2+0x25e - [un...
libLHAPDF.so!LHAPDF::Interpolator::interpolateXQ2+0x116 - [unknown sour...
libLHAPDF.so!LHAPDF::GridPDF::_xfxQ2+0x52 - [unknown source file]
libLHAPDF.so!LHAPDF::PDF::xfxQ2+0x9c - [unknown source file]
libLHAPDFSherpa.so!PDF::LHAPDF_CPP_Interface::GetXPDF+0x35a - LHAP...
libMCatNLOShowers.so.0!MCATNLO::Splitting_Function_Base::GetXPDF+0x1...
libMCatNLOShowers.so.0!MCATNLO::SF_Lorentz::JFI+0x75 - Splitting_Functi...
libMCatNLOShowers.so.0!MCATNLO::Shower::Reweight+0x2d0 - Shower.C:3...
libMCatNLOShowers.so.0!SHERPA::Variation_Weights::UpdateOrInitialiseWei...
libMCatNLOShowers.so.0!MCATNLO::Shower::EvolveShower+0xb5 - Shower....
libMCatNLOMain.so.0!MCATNLO::CS_MCatNLO::PerformMCatNLO+0xb3f - ...
libMCatNLOMain.so.0!MCATNLO::CS_MCatNLO::GeneratePoint+0x421 - CS...
libPhasicProcess.so.0!PHASIC::MCatNLO_Process::OneSEvent+0x28a - MCat...
libPhasicProcess.so.0!PHASIC::MCatNLO_Process::OneEvent+0x36f - MCatN...
libSherpaPerturbativePhysics.so.0!SHERPA::Matrix_Element_Handler::Generat...
libSherpaSingleEvents.so.0!SHERPA::Signal_Processes::Treat+0x60e - Signal_P...
libSherpaSingleEvents.so.0!SHERPA::Event_Handler::IterateEventPhases+0x3d...
libSherpaSingleEvents.so.0!SHERPA::Event_Handler::GenerateStandardPertur...
libSherpaSingleEvents.so.0!SHERPA::Event_Handler::GenerateEvent+0x590 - ...
libSherpaMain.so.0!SHERPA::Sherpa::GenerateOneEvent+0x28b - Sherpa.C:222
Sherpa!main+0xa7 - Main.C:27
libc.so.6!__libc_start_main+0xf4 - [unknown source file]
Sherpa!_start+0x28 - [unknown source file]

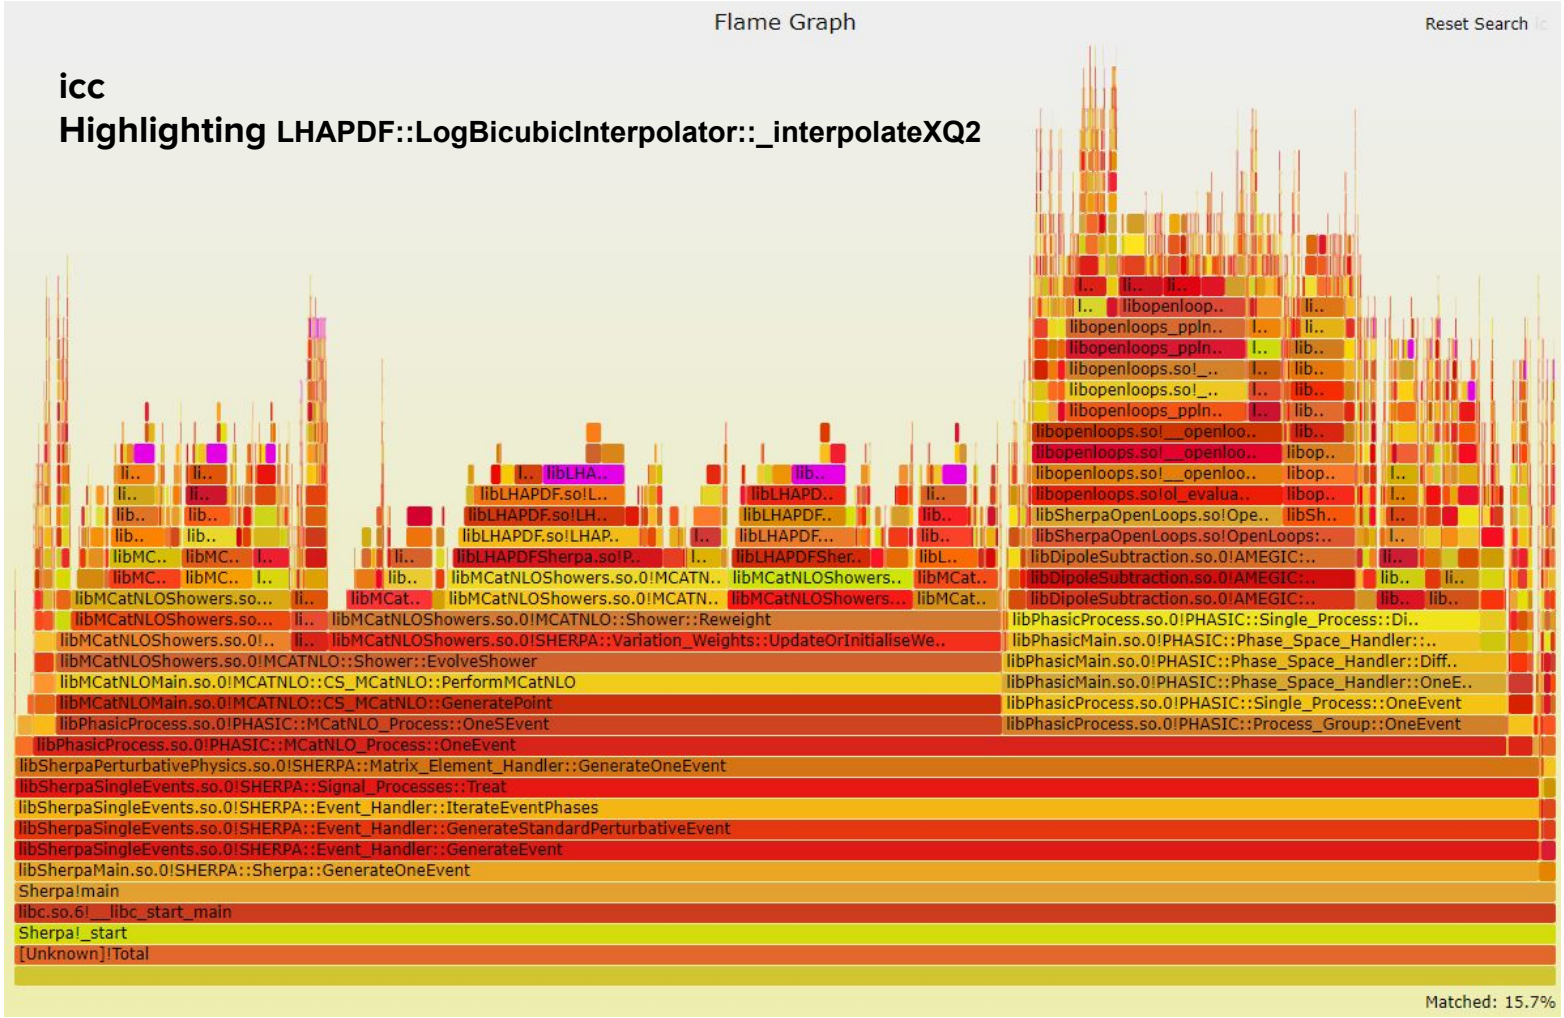
Flame Graph
Reset Search
icc
Highlighting LHAPDF::LogBicubicInterpolator::_interpolateXQ2
libMCatNLOShowers.so.0!MCATNLO::Shower::Reweight
libMCatNLOShowers.so.0!MCATNLO::Shower::EvolveShower
libMCatNLOMain.so.0!MCATNLO::CS_MCatNLO::PerformMCatNLO
libMCatNLOMain.so.0!MCATNLO::CS_MCatNLO::GeneratePoint
libPhasicProcess.so.0!PHASIC::MCatNLO_Process::OneSEvent
libPhasicProcess.so.0!PHASIC::MCatNLO_Process::OneEvent
libSherpaPerturbativePhysics.so.0!SHERPA::Matrix_Element_Handler::GenerateOneEvent
libSherpaSingleEvents.so.0!SHERPA::Signal_Processes::Treat
libSherpaSingleEvents.so.0!SHERPA::Event_Handler::IterateEventPhases
libSherpaSingleEvents.so.0!SHERPA::Event_Handler::GenerateStandardPerturbativeEvent
libSherpaSingleEvents.so.0!SHERPA::Event_Handler::GenerateEvent
libSherpaMain.so.0!SHERPA::Sherpa::GenerateOneEvent
Sherpa!main
libc.so.6!__libc_start_main
Sherpa!_start
[Unknown]!Total
libPhasicProcess.so.0!PHASIC::Single_Process::Di..
libPhasicMain.so.0!PHASIC::Phase_Space_Handler::..
libPhasicMain.so.0!PHASIC::Phase_Space_Handler::Diff..
libPhasicMain.so.0!PHASIC::Phase_Space_Handler::OneE..
libPhasicProcess.so.0!PHASIC::Single_Process::OneEvent
libPhasicProcess.so.0!PHASIC::Process_Group::OneEvent
Matched: 15.7%

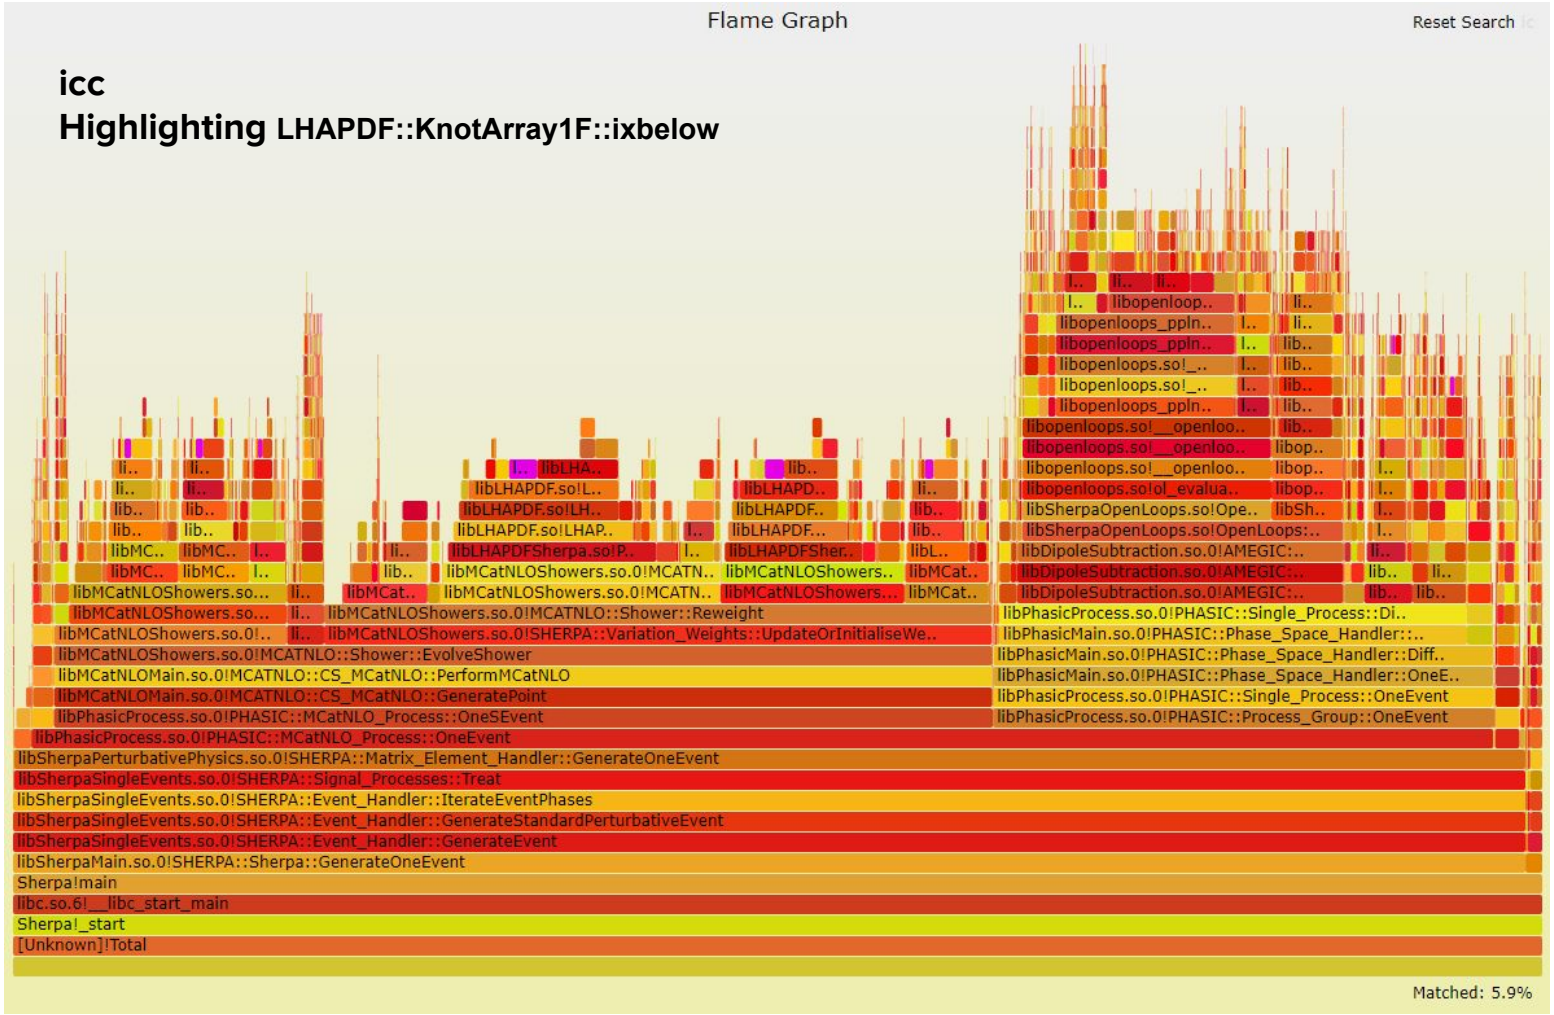
Flame Graph
Reset Search
icc
Highlighting LHAPDF::KnotArray1F::ixbelow
libLHAPDF.so!L..
libLHAPDF.so!LH..
libLHAPDF.so!LHAP..
libLHAPDFSherpa.so!P..
libMCatNLOShowers.so.0!MCATN..
libMCatNLOShowers.so.0!MCATNLO::Shower::Reweight
libMCatNLOShowers.so.0!SHERPA::Variation_Weights::UpdateOrInitialiseWe..
libMCatNLOShowers.so.0!MCATNLO::Shower::EvolveShower
libMCatNLOMain.so.0!MCATNLO::CS_MCatNLO::PerformMCatNLO
libMCatNLOMain.so.0!MCATNLO::CS_MCatNLO::GeneratePoint
libPhasicProcess.so.0!PHASIC::MCatNLO_Process::OneSEvent
libPhasicProcess.so.0!PHASIC::MCatNLO_Process::OneEvent
libSherpaPerturbativePhysics.so.0!SHERPA::Matrix_Element_Handler::GenerateOneEvent
libSherpaSingleEvents.so.0!SHERPA::Signal_Processes::Treat
libSherpaSingleEvents.so.0!SHERPA::Event_Handler::IterateEventPhases
libSherpaSingleEvents.so.0!SHERPA::Event_Handler::GenerateStandardPerturbativeEvent
libSherpaSingleEvents.so.0!SHERPA::Event_Handler::GenerateEvent
libSherpaMain.so.0!SHERPA::Sherpa::GenerateOneEvent
Sherpa!main
libc.so.6!__libc_start_main
Sherpa!_start
[Unknown]!Total
libopenloops.so!__openloo..
libopenloops.so!ol_evalua..
libSherpaOpenLoops.so!Ope..
libSherpaOpenLoops.so!OpenLoops:..
libDipoleSubtraction.so.0!AMEGIC:..
libPhasicProcess.so.0!PHASIC::Single_Process::Di..
libPhasicMain.so.0!PHASIC::Phase_Space_Handler::..
libPhasicMain.so.0!PHASIC::Phase_Space_Handler::Diff..
libPhasicMain.so.0!PHASIC::Phase_Space_Handler::OneE..
libPhasicProcess.so.0!PHASIC::Single_Process::OneEvent
libPhasicProcess.so.0!PHASIC::Process_Group::OneEvent
Matched: 5.9%

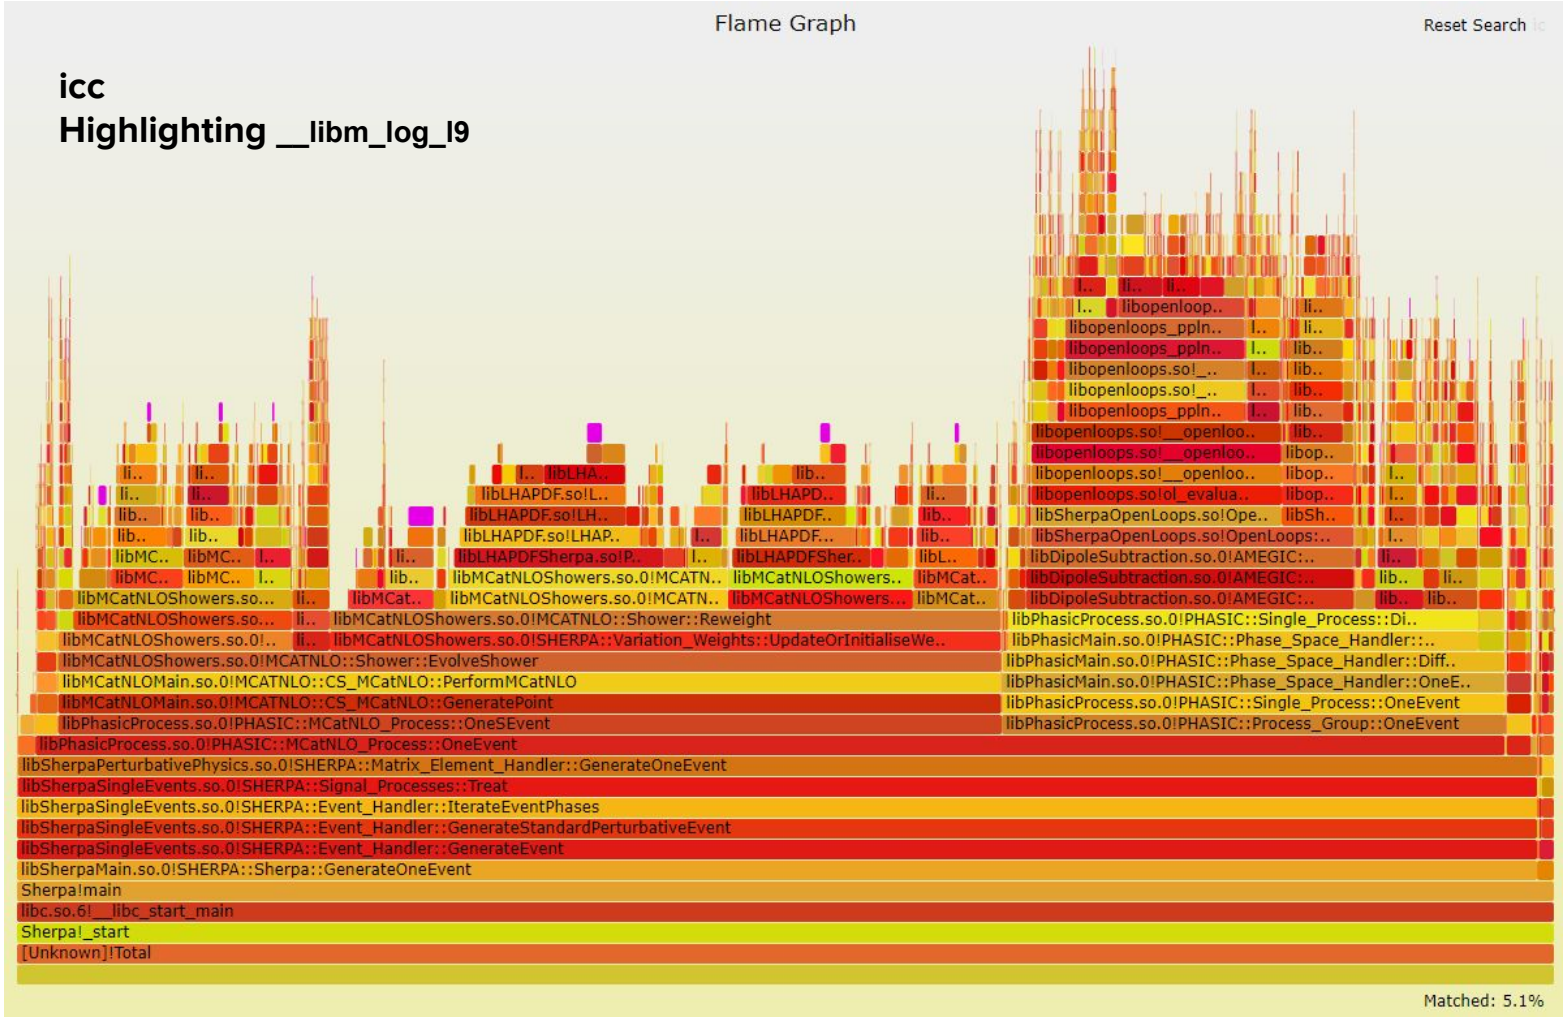
Flame Graph
Reset Search
icc
Highlighting __libm_log_l9
libMCatNLOShowers.so.0!MCATNLO::Shower::Reweight
libMCatNLOShowers.so.0!SHERPA::Variation_Weights::UpdateOrInitialiseWe..
libMCatNLOShowers.so.0!MCATNLO::Shower::EvolveShower
libMCatNLOMain.so.0!MCATNLO::CS_MCatNLO::PerformMCatNLO
libMCatNLOMain.so.0!MCATNLO::CS_MCatNLO::GeneratePoint
libPhasicProcess.so.0!PHASIC::MCatNLO_Process::OneSEvent
libPhasicProcess.so.0!PHASIC::MCatNLO_Process::OneEvent
libSherpaPerturbativePhysics.so.0!SHERPA::Matrix_Element_Handler::GenerateOneEvent
libSherpaSingleEvents.so.0!SHERPA::Signal_Processes::Treat
libSherpaSingleEvents.so.0!SHERPA::Event_Handler::IterateEventPhases
libSherpaSingleEvents.so.0!SHERPA::Event_Handler::GenerateStandardPerturbativeEvent
libSherpaSingleEvents.so.0!SHERPA::Event_Handler::GenerateEvent
libSherpaMain.so.0!SHERPA::Sherpa::GenerateOneEvent
Sherpa!main
libc.so.6!__libc_start_main
Sherpa!_start
[Unknown]!Total
libPhasicProcess.so.0!PHASIC::Single_Process::Di..
libPhasicMain.so.0!PHASIC::Phase_Space_Handler::..
libPhasicMain.so.0!PHASIC::Phase_Space_Handler::Diff..
libPhasicMain.so.0!PHASIC::Phase_Space_Handler::OneE..
libPhasicProcess.so.0!PHASIC::Single_Process::OneEvent
libPhasicProcess.so.0!PHASIC::Process_Group::OneEvent
Matched: 5.1%

Intel optmisations

- Intel have their own suite of optimisation flags
- Took a little effort to find a combination which would not cause compile errors in either LHAPDF or Sherpa
  - Sherpa: **`-xSSE4.2 -O3 -no-prec-div -fp-model fast=2`**
  - LHAPDF: **`-xSSE4.2 -O3`**
- Not currently using link time optimisation (**`-flto`**), issues...
- Total **18,598 s**
  - **6% faster than GCC**