Speed Improvements in Hit Merging

Doug Roberts

University of Maryland

Overview

- Prior to the introduction of backgrounds in FastSim, the Hit Merging code ran at a reasonable speed relative to other code
- However, with backgrounds the time spent in hit merging jumped way up.
- Reminder: the merging is basically a double-nested loop over SimHits. The more SimHits, the more time spent.
- Ran tests using:
  - V0.2.2
  - PacMCApp
  - `callgrind`
- Background mix used for tests was Bhabha, RadBhabha, Pair

Reconstruct

Simulate

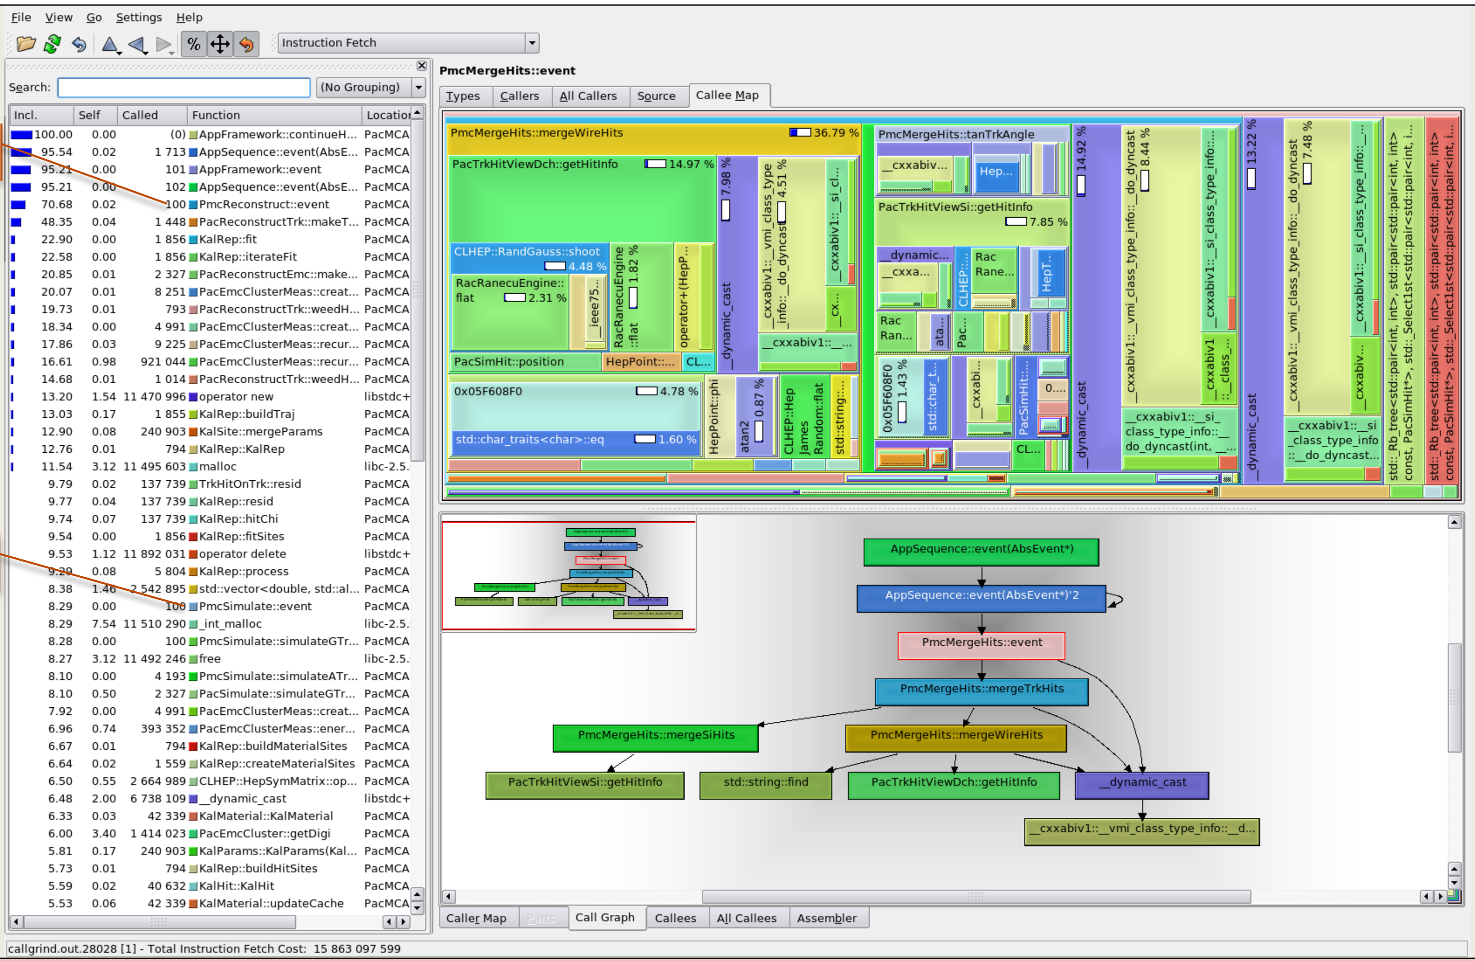
File View Go Settings Help
Instruction Fetch
Search:
(No Grouping)
PmcMergeHits::event
Types Callers All Callers Source Callee Map
PmcMergeHits::mergeWireHits 36.79 %
PacTrkHitViewDch::getHitInfo 14.97 %
PmcMergeHits::tanTrkAngle
PacTrkHitViewSi::getHitInfo 7.85 %
CLHEP::RandGauss::shoot 4.48 %
RacRanecuEngine::flat 2.31 %
PacSimHit::position
HepPoint::...
0x05F608F0 4.78 %
std::char_traits<char>::eq 1.60 %
__dynamic_cast
AppSequence::event(AbsEvent*)
AppSequence::event(AbsEvent*)'2
PmcMergeHits::event
PmcMergeHits::mergeTrkHits
PmcMergeHits::mergeSiHits
PmcMergeHits::mergeWireHits
PacTrkHitViewSi::getHitInfo
std::string::find
PacTrkHitViewDch::getHitInfo
__dynamic_cast
__cxxabiv1::__vmi_class_type_info::__d...
Caller Map Parts Call Graph Callees All Callees Assembler

Incl.	Self	Called	Function	Location
100.00	0.00	(0)	AppFramework::continueH...	PacMCA
95.54	0.02	1 713	AppSequence::event(AbsE...	PacMCA
95.21	0.00	101	AppFramework::event	PacMCA
95.21	0.00	102	AppSequence::event(AbsE...	PacMCA
70.68	0.02	100	PmcReconstruct::event	PacMCA
48.35	0.04	1 448	PacReconstructTrk::makeT...	PacMCA
22.90	0.00	1 856	KalRep::fit	PacMCA
22.58	0.00	1 856	KalRep::iterateFit	PacMCA
20.85	0.01	2 327	PacReconstructEmc::make...	PacMCA
20.07	0.01	8 251	PacEmcClusterMeas::creat...	PacMCA
19.73	0.01	793	PacReconstructTrk::weedH...	PacMCA
18.34	0.00	4 991	PacEmcClusterMeas::creat...	PacMCA
17.86	0.03	9 225	PacEmcClusterMeas::recur...	PacMCA
16.61	0.98	921 044	PacEmcClusterMeas::recur...	PacMCA
14.68	0.01	1 014	PacReconstructTrk::weedH...	PacMCA
13.20	1.54	11 470 996	operator new	libstdc+
13.03	0.17	1 855	KalRep::buildTraj	PacMCA
12.90	0.08	240 903	KalSite::mergeParams	PacMCA
12.76	0.01	794	KalRep::KalRep	PacMCA
11.54	3.12	11 495 603	malloc	libc-2.5.
9.79	0.02	137 739	TrkHitOnTrk::resid	PacMCA
9.77	0.04	137 739	KalRep::resid	PacMCA
9.74	0.07	137 739	KalRep::hitChi	PacMCA
9.54	0.00	1 856	KalRep::fitSites	PacMCA
9.53	1.12	11 892 031	operator delete	libstdc+
9.29	0.08	5 804	KalRep::process	PacMCA
8.38	1.46	2 542 895	std::vector<double, std::al...	PacMCA
8.29	0.00	100	PmcSimulate::event	PacMCA
8.29	7.54	11 510 290	_int_malloc	libc-2.5.
8.28	0.00	100	PmcSimulate::simulateGTr...	PacMCA
8.27	3.12	11 492 246	free	libc-2.5.
8.10	0.00	4 193	PmcSimulate::simulateATr...	PacMCA
8.10	0.50	2 327	PacSimulate::simulateGTr...	PacMCA
7.92	0.00	4 991	PacEmcClusterMeas::creat...	PacMCA
6.96	0.74	393 352	PacEmcClusterMeas::ener...	PacMCA
6.67	0.01	794	KalRep::buildMaterialSites	PacMCA
6.64	0.02	1 559	KalRep::createMaterialSites	PacMCA
6.50	0.55	2 664 989	CLHEP::HepSymMatrix::op...	PacMCA
6.48	2.00	6 738 109	__dynamic_cast	libstdc+
6.33	0.03	42 339	KalMaterial::KalMaterial	PacMCA
6.00	3.40	1 414 023	PacEmcCluster::getDigi	PacMCA
5.81	0.17	240 903	KalParams::KalParams(Kal...	PacMCA
5.73	0.01	794	KalRep::buildHitSites	PacMCA
5.59	0.02	40 632	KalHit::KalHit	PacMCA
5.53	0.06	42 339	KalMaterial::updateCache	PacMCA

callgrind.out.28028 [1] - Total Instruction Fetch Cost: 15 863 097 599

V0.2.2 Out-of-the-Box, no Background

PmcMergeHits::event doesn't appear above; it's too far down the list

Total Cost: 803,531,975 (5.07%) (not sure what the units are?)

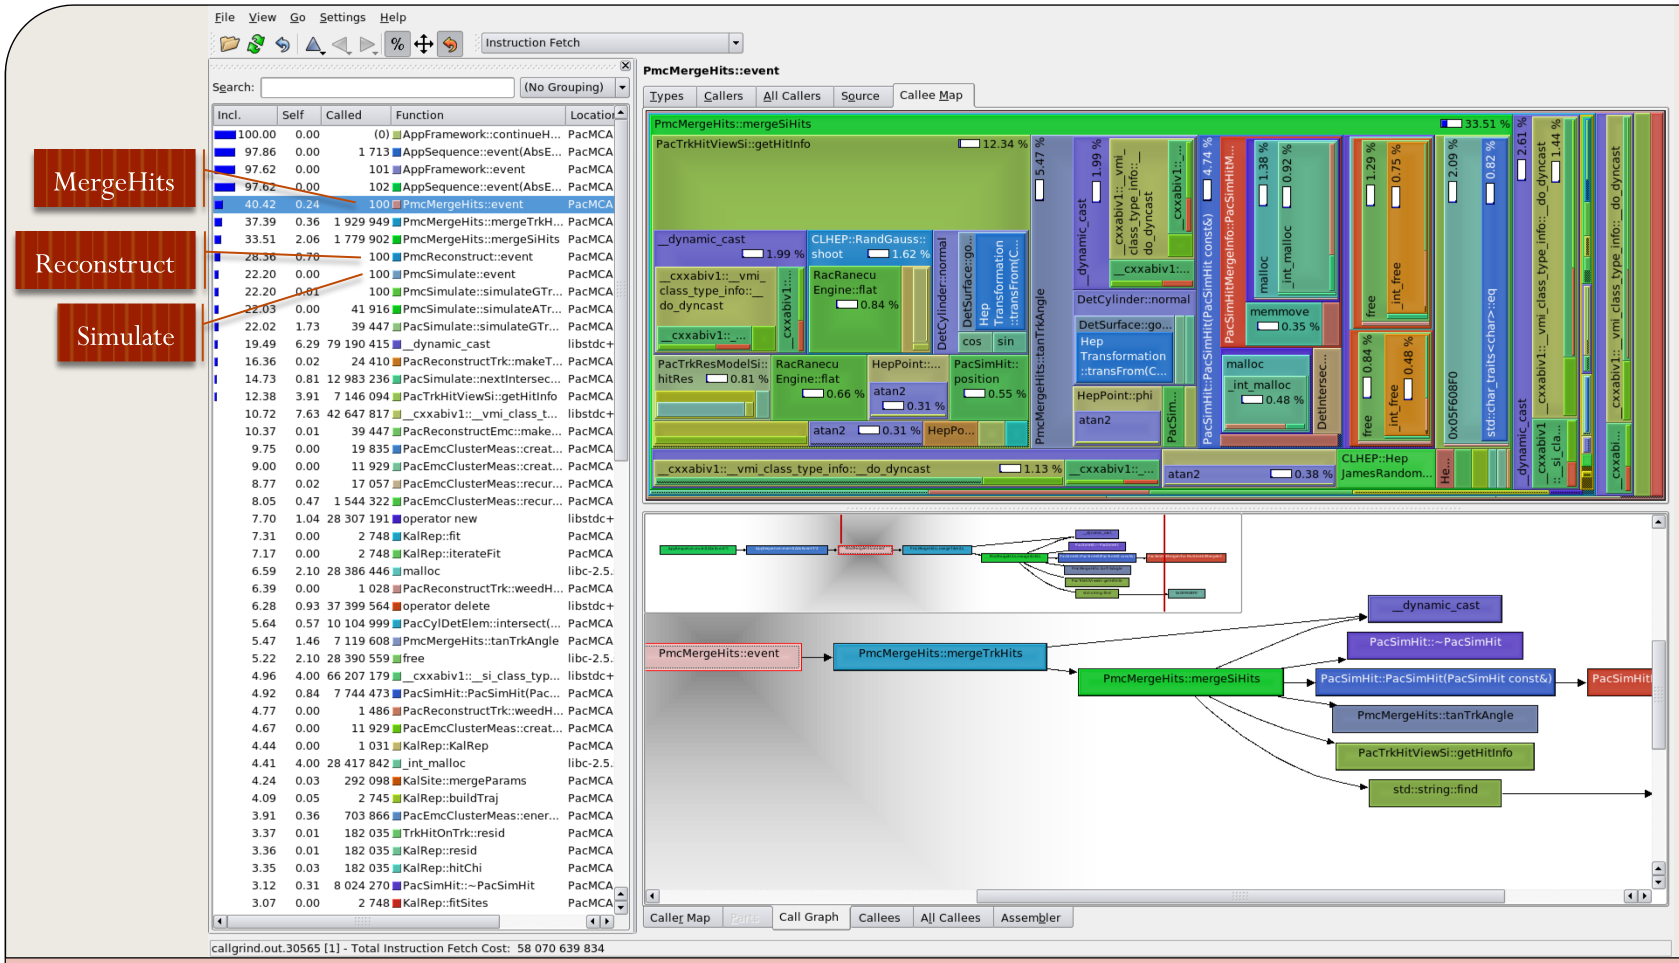
MergeHits
Reconstruct
Simulate
File
View
Go
Settings
Help
Instruction Fetch
Search:
(No Grouping)
PmcMergeHits::event
Types
Callers
All Callers
Source
Callee Map
PmcMergeHits::mergeSiHits 33.51 %
PacTrkHitViewSi::getHitInfo 12.34 %
__dynamic_cast 1.99 %
CLHEP::RandGauss:: shoot 1.62 %
RacRanecu Engine::flat 0.84 %
PacTrkResModelSi:: hitRes 0.81 %
RacRanecu Engine::flat 0.66 %
PacSimHit:: position 0.55 %
atan2 0.31 %
__cxxabiv1::__vmi_class_type_info::__do_dyncast 1.13 %
PmcMergeHits::tanTrkAngle 5.47 %
DetCylinder::normal
malloc
memmove 0.35 %
CLHEP::Hep JamesRandom...
PmcMergeHits::event
PmcMergeHits::mergeTrkHits
PmcMergeHits::mergeSiHits
__dynamic_cast
PacSimHit::~PacSimHit
PacSimHit::PacSimHit(PacSimHit const&)
PacSimHit
PmcMergeHits::tanTrkAngle
PacTrkHitViewSi::getHitInfo
std::string::find
Caller Map
Parts
Call Graph
Callees
All Callees
Assembler

Incl.	Self	Called	Function	Location
100.00	0.00	(0)	AppFramework::continueH...	PacMCA
97.86	0.00	1 713	AppSequence::event(AbsE...	PacMCA
97.62	0.00	101	AppFramework::event	PacMCA
97.62	0.00	102	AppSequence::event(AbsE...	PacMCA
40.42	0.24	100	PmcMergeHits::event	PacMCA
37.39	0.36	1 929 949	PmcMergeHits::mergeTrkH...	PacMCA
33.51	2.06	1 779 902	PmcMergeHits::mergeSiHits	PacMCA
28.36	0.70	100	PmcReconstruct::event	PacMCA
22.20	0.00	100	PmcSimulate::event	PacMCA
22.20	0.01	100	PmcSimulate::simulateGTr...	PacMCA
22.03	0.00	41 916	PmcSimulate::simulateATr...	PacMCA
22.02	1.73	39 447	PacSimulate::simulateGTr...	PacMCA
19.49	6.29	79 190 415	__dynamic_cast	libstdc+
16.36	0.02	24 410	PacReconstructTrk::makeT...	PacMCA
14.73	0.81	12 983 236	PacSimulate::nextIntersec...	PacMCA
12.38	3.91	7 146 094	PacTrkHitViewSi::getHitInfo	PacMCA
10.72	7.63	42 647 817	__cxxabiv1::__vmi_class_t...	libstdc+
10.37	0.01	39 447	PacReconstructEmc::make...	PacMCA
9.75	0.00	19 835	PacEmcClusterMeas::creat...	PacMCA
9.00	0.00	11 929	PacEmcClusterMeas::creat...	PacMCA
8.77	0.02	17 057	PacEmcClusterMeas::recur...	PacMCA
8.05	0.47	1 544 322	PacEmcClusterMeas::recur...	PacMCA
7.70	1.04	28 307 191	operator new	libstdc+
7.31	0.00	2 748	KalRep::fit	PacMCA
7.17	0.00	2 748	KalRep::iterateFit	PacMCA
6.59	2.10	28 386 446	malloc	libc-2.5.
6.39	0.00	1 028	PacReconstructTrk::weedH...	PacMCA
6.28	0.93	37 399 564	operator delete	libstdc+
5.64	0.57	10 104 999	PacCylDetElem::intersect(...	PacMCA
5.47	1.46	7 119 608	PmcMergeHits::tanTrkAngle	PacMCA
5.22	2.10	28 390 559	free	libc-2.5.
4.96	4.00	66 207 179	__cxxabiv1::__si_class_typ...	libstdc+
4.92	0.84	7 744 473	PacSimHit::PacSimHit(Pac...	PacMCA
4.77	0.00	1 486	PacReconstructTrk::weedH...	PacMCA
4.67	0.00	11 929	PacEmcClusterMeas::creat...	PacMCA
4.44	0.00	1 031	KalRep::KalRep	PacMCA
4.41	4.00	28 417 842	_int_malloc	libc-2.5.
4.24	0.03	292 098	KalSite::mergeParams	PacMCA
4.09	0.05	2 745	KalRep::buildTraj	PacMCA
3.91	0.36	703 866	PacEmcClusterMeas::ener...	PacMCA
3.37	0.01	182 035	TrkHitOnTrk::resid	PacMCA
3.36	0.01	182 035	KalRep::resid	PacMCA
3.35	0.03	182 035	KalRep::hitChi	PacMCA
3.12	0.31	8 024 270	PacSimHit::~PacSimHit	PacMCA
3.07	0.00	2 748	KalRep::fitSites	PacMCA

```
callgrind.out.30565 [1] - Total Instruction Fetch Cost:  58 070 639 834
```

V0.2.2 Out-of-the-Box with Background

PmcMergeHits jumps above PmcReconstruct and PmcSimulate! Becomes single most expensive module

Total Cost: 23,470,429,356 (40.42%)

Nearly a factor of 30 increase in cost

Speed Improvements

- Using `callgrind` output, it was pretty easy to see where the time was being spent
- Strategy:
  - Cache any expensive calculations
    - Created a new object to store the cached info
      - Easy. There was no dynamic data.
    - Includes caching results from some expensive `dynamic_cast` calls
  - Rearrange cuts: cheap early, expensive late
  - Found some functions that were passing objects as arguments instead of pointers
    - Lots of copy c'tor calls
  - Got rid of a `std::string::find()` call

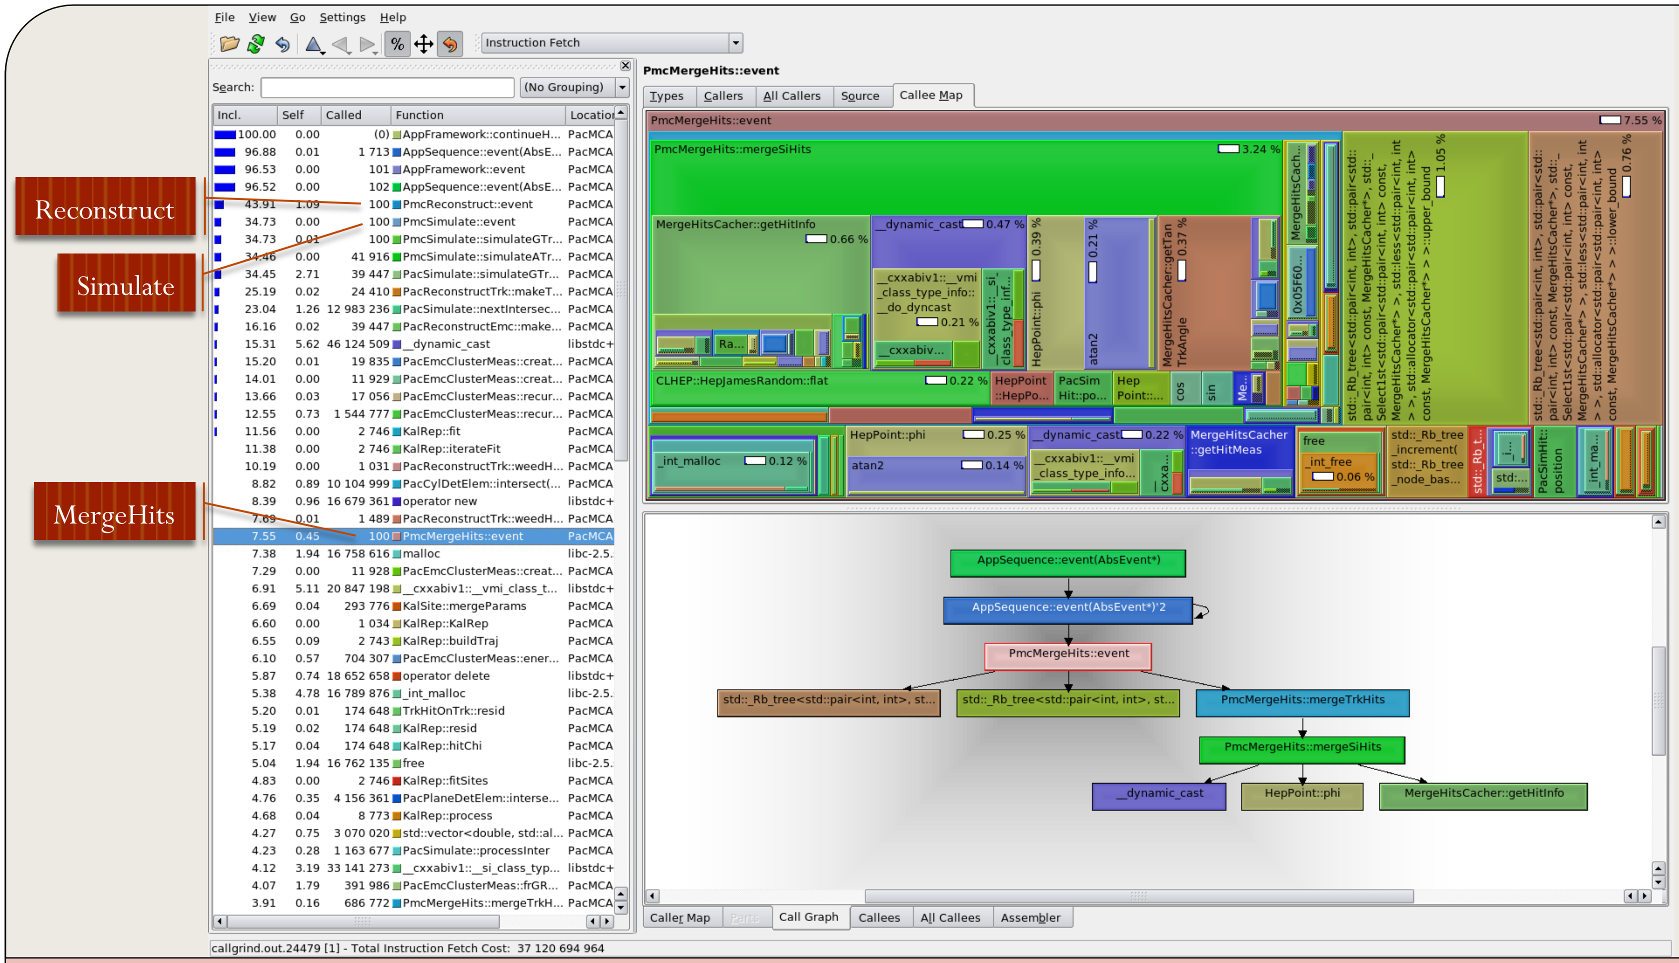
Reconstruct
Simulate
MergeHits
File View Go Settings Help
Instruction Fetch
Search:
(No Grouping)
PmcMergeHits::event
Types Callers All Callers Source Callee Map
PmcMergeHits::event 7.55 %
PmcMergeHits::mergeSiHits 3.24 %
MergeHitsCacher::getHitInfo 0.66 %
_dynamic_cast 0.47 %
CLHEP::HepJamesRandom::flat 0.22 %
HepPoint::phi 0.25 %
atan2 0.14 %
_int_malloc 0.12 %
MergeHitsCacher ::getHitMeas
free
_int_free 0.06 %
AppSequence::event(AbsEvent*)
AppSequence::event(AbsEvent*)'2
PmcMergeHits::event
std::_Rb_tree<std::pair<int, int>, st...
std::_Rb_tree<std::pair<int, int>, st...
PmcMergeHits::mergeTrkHits
PmcMergeHits::mergeSiHits
_dynamic_cast
HepPoint::phi
MergeHitsCacher::getHitInfo
Caller Map Parts Call Graph Callees All Callees Assembler

Incl.	Self	Called	Function	Location
100.00	0.00	(0)	AppFramework::continueH...	PacMCA
96.88	0.01	1 713	AppSequence::event(AbsE...	PacMCA
96.53	0.00	101	AppFramework::event	PacMCA
96.52	0.00	102	AppSequence::event(AbsE...	PacMCA
43.91	1.09	100	PmcReconstruct::event	PacMCA
34.73	0.00	100	PmcSimulate::event	PacMCA
34.73	0.01	100	PmcSimulate::simulateGTr...	PacMCA
34.46	0.00	41 916	PmcSimulate::simulateATr...	PacMCA
34.45	2.71	39 447	PacSimulate::simulateGTr...	PacMCA
25.19	0.02	24 410	PacReconstructTrk::makeT...	PacMCA
23.04	1.26	12 983 236	PacSimulate::nextIntersec...	PacMCA
16.16	0.02	39 447	PacReconstructEmc::make...	PacMCA
15.31	5.62	46 124 509	__dynamic_cast	libstdc+
15.20	0.01	19 835	PacEmcClusterMeas::creat...	PacMCA
14.01	0.00	11 929	PacEmcClusterMeas::creat...	PacMCA
13.66	0.03	17 056	PacEmcClusterMeas::recur...	PacMCA
12.55	0.73	1 544 777	PacEmcClusterMeas::recur...	PacMCA
11.56	0.00	2 746	KalRep::fit	PacMCA
11.38	0.00	2 746	KalRep::iterateFit	PacMCA
10.19	0.00	1 031	PacReconstructTrk::weedH...	PacMCA
8.82	0.89	10 104 999	PacCylDetElem::intersect(...	PacMCA
8.39	0.96	16 679 361	operator new	libstdc+
7.69	0.01	1 489	PacReconstructTrk::weedH...	PacMCA
7.55	0.45	100	PmcMergeHits::event	PacMCA
7.38	1.94	16 758 616	malloc	libc-2.5.
7.29	0.00	11 928	PacEmcClusterMeas::creat...	PacMCA
6.91	5.11	20 847 198	__cxxabiv1::__vmi_class_t...	libstdc+
6.69	0.04	293 776	KalSite::mergeParams	PacMCA
6.60	0.00	1 034	KalRep::KalRep	PacMCA
6.55	0.09	2 743	KalRep::buildTraj	PacMCA
6.10	0.57	704 307	PacEmcClusterMeas::ener...	PacMCA
5.87	0.74	18 652 658	operator delete	libstdc+
5.38	4.78	16 789 876	_int_malloc	libc-2.5.
5.20	0.01	174 648	TrkHitOnTrk::resid	PacMCA
5.19	0.02	174 648	KalRep::resid	PacMCA
5.17	0.04	174 648	KalRep::hitChi	PacMCA
5.04	1.94	16 762 135	free	libc-2.5.
4.83	0.00	2 746	KalRep::fitSites	PacMCA
4.76	0.35	4 156 361	PacPlaneDetElem::interse...	PacMCA
4.68	0.04	8 773	KalRep::process	PacMCA
4.27	0.75	3 070 020	std::vector<double, std::al...	PacMCA
4.23	0.28	1 163 677	PacSimulate::processInter	PacMCA
4.12	3.19	33 141 273	__cxxabiv1::__si_class_typ...	libstdc+
4.07	1.79	391 986	PacEmcClusterMeas::frGR...	PacMCA
3.91	0.16	686 772	PmcMergeHits::mergeTrkH...	PacMCA

```
callgrind.out.24479 [1] - Total Instruction Fetch Cost:  37 120 694 964
```

After speed fixes, with background

Total Cost: 2,804,375,274 (7.55%)

More than a factor of 8 improvement

Only ~3.5 times slower than original no-background version

For comparison, Simulate is ~10 times slower, Reconstruct is only ~1.5 times slower

Conclusion

- Able to get the merging code's speed to a (hopefully) acceptable level
- There should be virtually no difference in the output. It's just faster
- `callgrind` is a useful tool!