

GammaLib - Action #3203

Implement spatial IRF integration methods that return IRF values for all events

04/11/2020 06:57 PM - Knödlseder Jürgen

Status:	Closed	Start date:	04/11/2020
Priority:	Normal	Due date:	
Assigned To:	Knödlseder Jürgen	% Done:	100%
Category:		Estimated time:	0.00 hour
Target version:	2.0.0		
Description			
Implement a			
std::vector<double> irf(const GSource& source, const GObservation& obs) const;			
that returns a vector of model events for a given source model, where each element of the vector corresponds to one event. This should leads to a much more efficient computation of the spatially integrated model values, since all the coordinate transformation for a given model are done once. Since the irf computation is in the innermost loop over the integration (e.g. in GResponse::irf_radial_kern_phi::eval() for the radial model), one can easily loop there over all events to get the IRF for each event.			
To make this work the GObservation::likelihood_poisson_unbinned(), GObservation::likelihood_poisson_binned() and GObservation::likelihood_gaussian_binned() methods need to be modified to work on entire vectors of events, the same is needs to be done for the GObservation::model() and the GModel::eval() method that should operate directly on the entire observation.			
Related issues:			
Related to GammaLib - Feature # 3328: Implement numerical integration over ve...		Closed	09/03/2020

History

#1 - 04/13/2020 10:35 AM - Knödlseder Jürgen

- File gammalib_tn0003_response.pdf added
- Status changed from New to In Progress
- % Done changed from 0 to 10

I added the Tech Note attachment:gammalib_tn0003_response.pdf to describe a possible modified scheme.

#2 - 09/03/2020 05:21 PM - Knödlseder Jürgen

- Related to Feature #3328: Implement numerical integration over vectors of functions added

#3 - 09/11/2020 04:50 PM - Knödlseder Jürgen

- Assigned To set to Knödlseder Jürgen
- Target version set to 2.0.0

The following methods need to be developed:

```
virtual GNdarray GModels::eval(const GObservation& obs,
                               const bool& gradients = false) const;
virtual GNdarray GModel::eval(const GObservation& obs,
                             const bool& gradients = false) const = 0;
virtual GNdarray GModelSky::eval(const GObservation& obs,
                                 const bool& gradients) const
virtual GNdarray GResponse::convolve(const GModelSky& model,
                                     const GObservation& obs,
                                     const bool& grad) const
virtual GNdarray GResponse::irf(const GPhoton& photon,
                               const GObservation& obs) const = 0;
virtual GNdarray GResponse::irf_spatial(const GSource& source,
```

```
    const GObservation& obs) const;  
virtual GNdarray GObservation::model(const GModels& models) const
```

Note that the gradients for all events will be packed into GNdarray for GObservation::model which will return a 2D array, where one dimension are all events, and the second dimension are the model value followed by all model gradients. Eventually, a new class could be introduced that provides specific access methods to the information.

Once the methods exists, the model evaluation can be put outside the event loop in

```
double GObservation::likelihood_poisson_unbinned()  
double GObservation::likelihood_poisson_binned()  
double GObservation::likelihood_gaussian_binned()
```

#4 - 09/19/2020 10:08 AM - Knödlseder Jürgen

- % Done changed from 10 to 20

I implemented the methods

```
GVector GModel::eval(const GObservation& obs, GMMatrixSparse* gradients)  
GVector GModelSky::eval(const GObservation& obs, GMMatrixSparse* gradients)
```

that return GVector instances that contain the model values for each event. Optionally, if gradients != NULL, the methods return a matrix with the parameter gradients, where each column contains the gradients for a given event, and the rows contain the gradients for all parameters in the model container.

The GModelSky::eval() method makes use of the new method

```
GVector GResponse::convolve(const GModelSky& model, const GObservation& obs, GMMatrixSparse* gradients)
```

that convolves the sky model with the instrument response function.

The GObservation::likelihood_poisson_binned() and GObservation::likelihood_gaussian_binned() methods now call the

```
GVector GObservation::model(const GModels& model, GMMatrixSparse* gradients)
```

method that returns a GVector instance that contains the model value for each event. Optionally, if gradients != NULL, the method returns a matrix with the parameter gradients, where each row contains the gradients for a given event, and the columns contain the gradients for all parameters in the model container (note that the ordering is different to the one returned by the GModel::eval() and GModelSky::eval() methods).

I also implemented the method

```
GVector GObservation::model_grad(const GModel& model, const GModelPar& par)
```

that computes the numerical parameter gradient for a given model parameter for all events in a given observation.

The code was checked with respect to a reference run using the old methods. The same results were obtained.

The next step is to remove the event loop from the

```
GVector GResponse::convolve(const GModelSky& model, const GObservation& obs, GMMatrixSparse* gradients)
```

method and use vector functions.

#5 - 10/02/2020 11:50 AM - Knödlseder Jürgen

I fully implemented the vector integration for radial models in a stacked analysis. To test the code, I generated a simulation of the Crab with a spatial Gaussian model. Here is the clike output for the fit using the old event wise treatment:

```
2020-09-15T14:07:09: +=====
2020-09-15T14:07:09: | Maximum likelihood optimisation |
2020-09-15T14:07:09: +=====
2020-09-15T14:08:26: >Iteration 0: -logL=-118103.118, Lambda=1.0e-03
2020-09-15T14:09:43: >Iteration 1: -logL=-118106.615, Lambda=1.0e-03, delta=3.497, step=1.0e+00, max(|grad|)=19.603365 [Sigma:2]
2020-09-15T14:10:59: >Iteration 2: -logL=-118106.615, Lambda=1.0e-04, delta=0.000, step=1.0e+00, max(|grad|)=0.215447 [Sigma:2]
2020-09-15T14:12:17:
2020-09-15T14:12:17: +=====
2020-09-15T14:12:17: | Maximum likelihood optimisation results |
2020-09-15T14:12:17: +=====
2020-09-15T14:12:17: === GOptimizerLM ===
2020-09-15T14:12:17: Optimized function value ...: -118106.615
2020-09-15T14:12:17: Absolute precision .....: 0.005
2020-09-15T14:12:17: Acceptable value decrease ..: 2
2020-09-15T14:12:17: Optimization status .....: converged
2020-09-15T14:12:17: Number of parameters .....: 11
2020-09-15T14:12:17: Number of free parameters ..: 7
2020-09-15T14:12:17: Number of iterations .....: 2
2020-09-15T14:12:17: Lambda .....: 1e-05
2020-09-15T14:12:17: Maximum log likelihood ....: 118106.615
2020-09-15T14:12:17: Observed events (Nobs) ....: 784404.000
2020-09-15T14:12:17: Predicted events (Npred) ....: 784404.001 (Nobs - Npred = -0.000732269138097763)
2020-09-15T14:12:17: === GModels ===
2020-09-15T14:12:17: Number of models .....: 2
2020-09-15T14:12:17: Number of parameters .....: 11
2020-09-15T14:12:17: === GModelSky ===
2020-09-15T14:12:17: Name .....: Crab
2020-09-15T14:12:17: Instruments .....: all
2020-09-15T14:12:17: Observation identifiers ...: all
2020-09-15T14:12:17: Model type .....: ExtendedSource
2020-09-15T14:12:17: Model components .....: "RadialGaussian" * "PowerLaw" * "Constant"
2020-09-15T14:12:17: Number of parameters .....: 7
2020-09-15T14:12:17: Number of spatial par's ...: 3
2020-09-15T14:12:17: RA .....: 83.6322984937305 +/- 0.0013688053300792 [-360,360] deg (free,scale=1)
2020-09-15T14:12:17: DEC .....: 22.0132602065802 +/- 0.0012665661852356 [-90,90] deg (free,scale=1)
2020-09-15T14:12:17: Sigma .....: 0.199377918253211 +/- 0.000855412014669404 [0.01,10] deg (free,scale=1)
2020-09-15T14:12:17: Number of spectral par's ...: 3
2020-09-15T14:12:17: Prefactor .....: 5.65581461097373e-16 +/- 3.33104053091683e-18 [1e-23,1e-13] ph/cm2/s/MeV
(free,scale=1e-16,gradient)
2020-09-15T14:12:17: Index .....: -2.4738735499866 +/- 0.0046302759481871 [-5,-0] (free,scale=-1,gradient)
2020-09-15T14:12:17: PivotEnergy .....: 300000 [10000,1000000000] MeV (fixed,scale=1000000,gradient)
2020-09-15T14:12:17: Number of temporal par's ..: 1
2020-09-15T14:12:17: Normalization .....: 1 (relative value) (fixed,scale=1,gradient)
2020-09-15T14:12:17: Number of scale par's ....: 0
2020-09-15T14:12:17: === GCTAModelCubeBackground ===
2020-09-15T14:12:17: Name .....: BackgroundModel
2020-09-15T14:12:17: Instruments .....: CTA, HESS, MAGIC, VERITAS
2020-09-15T14:12:17: Observation identifiers ...: all
2020-09-15T14:12:17: Model type .....: "PowerLaw" * "Constant"
2020-09-15T14:12:17: Number of parameters .....: 4
2020-09-15T14:12:17: Number of spectral par's ...: 3
```

```

2020-09-15T14:12:17: Prefactor .....: 0.997217425824902 +/- 0.00234232407761604 [0.01,100] ph/cm2/s/MeV (free,scale=1,gradient)
2020-09-15T14:12:17: Index .....: -0.00232209919585069 +/- 0.00134153129520882 [-5,5] (free,scale=1,gradient)
2020-09-15T14:12:17: PivotEnergy .....: 1000000 MeV (fixed,scale=1000000,gradient)
2020-09-15T14:12:17: Number of temporal par's ...: 1
2020-09-15T14:12:17: Normalization .....: 1 (relative value) (fixed,scale=1,gradient)
2020-09-15T14:12:17:
2020-09-15T14:12:17: +=====
2020-09-15T14:12:17: | Save results |
2020-09-15T14:12:17: +=====
2020-09-15T14:12:17: Model definition file ....: crab_results.xml
2020-09-15T14:12:17: Covariance matrix file ....: NONE
2020-09-15T14:12:17:
2020-09-15T14:12:17: Application "ctlike" terminated after 332 wall clock seconds, consuming 307.796 seconds of CPU time.

```

Using the vector integration, the following results are achieved (the speed-up is about of a factor of 4.5):

```

2020-09-22T20:30:08: +=====
2020-09-22T20:30:08: | Maximum likelihood optimisation |
2020-09-22T20:30:08: +=====
2020-09-22T20:30:25: >Iteration 0: -logL=-118103.118, Lambda=1.0e-03
2020-09-22T20:30:42: >Iteration 1: -logL=-118106.615, Lambda=1.0e-03, delta=3.497, step=1.0e+00, max(|grad|)=19.603365 [Sigma:2]
2020-09-22T20:30:59: >Iteration 2: -logL=-118106.615, Lambda=1.0e-04, delta=0.000, step=1.0e+00, max(|grad|)=0.215447 [Sigma:2]
2020-09-22T20:31:16:
2020-09-22T20:31:16: +=====
2020-09-22T20:31:16: | Maximum likelihood optimisation results |
2020-09-22T20:31:16: +=====
2020-09-22T20:31:16: === GOptimizerLM ===
2020-09-22T20:31:16: Optimized function value ...: -118106.615
2020-09-22T20:31:16: Absolute precision .....: 0.005
2020-09-22T20:31:16: Acceptable value decrease ..: 2
2020-09-22T20:31:16: Optimization status .....: converged
2020-09-22T20:31:16: Number of parameters .....: 11
2020-09-22T20:31:16: Number of free parameters ..: 7
2020-09-22T20:31:16: Number of iterations .....: 2
2020-09-22T20:31:16: Lambda .....: 1e-05
2020-09-22T20:31:16: Maximum log likelihood ....: 118106.615
2020-09-22T20:31:16: Observed events (Nobs) ...: 784404.000
2020-09-22T20:31:16: Predicted events (Npred) ...: 784404.001 (Nobs - Npred = -0.00073208415415138)
2020-09-22T20:31:16: === GModels ===
2020-09-22T20:31:16: Number of models .....: 2
2020-09-22T20:31:16: Number of parameters .....: 11
2020-09-22T20:31:16: === GModelSky ===
2020-09-22T20:31:16: Name .....: Crab
2020-09-22T20:31:16: Instruments .....: all
2020-09-22T20:31:16: Observation identifiers ...: all
2020-09-22T20:31:16: Model type .....: ExtendedSource
2020-09-22T20:31:16: Model components .....: "RadialGaussian" * "PowerLaw" * "Constant"
2020-09-22T20:31:16: Number of parameters .....: 7
2020-09-22T20:31:16: Number of spatial par's ...: 3
2020-09-22T20:31:16: RA .....: 83.6322984937299 +/- 0.00136880533017678 [-360,360] deg (free,scale=1)
2020-09-22T20:31:16: DEC .....: 22.013260206577 +/- 0.00126656618514527 [-90,90] deg (free,scale=1)
2020-09-22T20:31:16: Sigma .....: 0.19937791825324 +/- 0.000855412014670356 [0.01,10] deg (free,scale=1)
2020-09-22T20:31:16: Number of spectral par's ...: 3
2020-09-22T20:31:16: Prefactor .....: 5.65581461096391e-16 +/- 3.3310405309414e-18 [1e-23,1e-13] ph/cm2/s/MeV (free,scale=1e-16,gradient)
2020-09-22T20:31:16: Index .....: -2.4738735499867 +/- 0.0046302759481897 [-5,-0] (free,scale=-1,gradient)
2020-09-22T20:31:16: PivotEnergy .....: 300000 [10000,1000000000] MeV (fixed,scale=1000000,gradient)
2020-09-22T20:31:16: Number of temporal par's ...: 1
2020-09-22T20:31:16: Normalization .....: 1 (relative value) (fixed,scale=1,gradient)
2020-09-22T20:31:16: Number of scale par's .....: 0
2020-09-22T20:31:16: === GCTAModelCubeBackground ===
2020-09-22T20:31:16: Name .....: BackgroundModel
2020-09-22T20:31:16: Instruments .....: CTA, HESS, MAGIC, VERITAS
2020-09-22T20:31:16: Observation identifiers ...: all
2020-09-22T20:31:16: Model type .....: "PowerLaw" * "Constant"
2020-09-22T20:31:16: Number of parameters .....: 4
2020-09-22T20:31:16: Number of spectral par's ...: 3
2020-09-22T20:31:16: Prefactor .....: 0.997217425824893 +/- 0.00234232407761589 [0.01,100] ph/cm2/s/MeV (free,scale=1,gradient)
2020-09-22T20:31:16: Index .....: -0.00232209919585454 +/- 0.00134153129520875 [-5,5] (free,scale=1,gradient)
2020-09-22T20:31:16: PivotEnergy .....: 1000000 MeV (fixed,scale=1000000,gradient)
2020-09-22T20:31:16: Number of temporal par's ...: 1
2020-09-22T20:31:16: Normalization .....: 1 (relative value) (fixed,scale=1,gradient)
2020-09-22T20:31:16:

```

```

2020-09-22T20:31:16: +=====
2020-09-22T20:31:16: | Save results |
2020-09-22T20:31:16: +=====
2020-09-22T20:31:16: Model definition file .....: crab_results.xml
2020-09-22T20:31:16: Covariance matrix file ....: NONE
2020-09-22T20:31:16:
2020-09-22T20:31:16: Application "ctlike" terminated after 68 wall clock seconds, consuming 67.9258 seconds of CPU time.

```

#6 - 10/02/2020 11:54 AM - Knödlseder Jürgen

I finally managed to implement also the computation of analytical spatial gradients for the radial Gaussian model using stacked analysis. The results are shown below. The speed-up is now a factor of 6.7 with respect to the reference analysis:

```

2020-10-02T09:41:03: +=====
2020-10-02T09:41:03: | Maximum likelihood optimisation |
2020-10-02T09:41:03: +=====
2020-10-02T09:41:14: >Iteration 0: -logL=-118103.118, Lambda=1.0e-03
2020-10-02T09:41:25: >Iteration 1: -logL=-118106.615, Lambda=1.0e-03, delta=3.497, step=1.0e+00, max(|grad|)=19.142435 [Sigma:2]
2020-10-02T09:41:37: >Iteration 2: -logL=-118106.615, Lambda=1.0e-04, delta=0.000, step=1.0e+00, max(|grad|)=0.204245 [Sigma:2]
2020-10-02T09:41:48:
2020-10-02T09:41:48: +=====
2020-10-02T09:41:48: | Maximum likelihood optimisation results |
2020-10-02T09:41:48: +=====
2020-10-02T09:41:48: === GOptimizerLM ===
2020-10-02T09:41:48: Optimized function value ...: -118106.615
2020-10-02T09:41:48: Absolute precision .....: 0.005
2020-10-02T09:41:48: Acceptable value decrease ..: 2
2020-10-02T09:41:48: Optimization status .....: converged
2020-10-02T09:41:48: Number of parameters .....: 11
2020-10-02T09:41:48: Number of free parameters ..: 7
2020-10-02T09:41:48: Number of iterations .....: 2
2020-10-02T09:41:48: Lambda .....: 1e-05
2020-10-02T09:41:48: Maximum log likelihood .....: 118106.615
2020-10-02T09:41:48: Observed events (Nobs) ....: 784404.000
2020-10-02T09:41:48: Predicted events (Npred) ....: 784404.001 (Nobs - Npred = -0.000611254596151412)
2020-10-02T09:41:48: === GModels ===
2020-10-02T09:41:48: Number of models .....: 2
2020-10-02T09:41:48: Number of parameters .....: 11
2020-10-02T09:41:48: === GModelSky ===
2020-10-02T09:41:48: Name .....: Crab
2020-10-02T09:41:48: Instruments .....: all
2020-10-02T09:41:48: Observation identifiers ...: all
2020-10-02T09:41:48: Model type .....: ExtendedSource
2020-10-02T09:41:48: Model components .....: "RadialGaussian" * "PowerLaw" * "Constant"
2020-10-02T09:41:48: Number of parameters .....: 7
2020-10-02T09:41:48: Number of spatial par's ...: 3
2020-10-02T09:41:48: RA .....: 83.6322984504952 +/- 0.00136879926594665 [-360,360] deg (free,scale=1)
2020-10-02T09:41:48: DEC .....: 22.0132604173708 +/- 0.0012665604200981 [-90,90] deg (free,scale=1)
2020-10-02T09:41:48: Sigma .....: 0.199377214322206 +/- 0.000855400486036962 [0.01,10] deg (free,scale=1,gradient)
2020-10-02T09:41:48: Number of spectral par's ...: 3
2020-10-02T09:41:48: Prefactor .....: 5.65580336702951e-16 +/- 3.33102267159428e-18 [1e-23,1e-13] ph/cm2/s/MeV
(free,scale=1e-16,gradient)
2020-10-02T09:41:48: Index .....: -2.47387340817538 +/- 0.00463027895322123 [-5,-0] (free,scale=-1,gradient)
2020-10-02T09:41:48: PivotEnergy .....: 300000 [10000,1000000000] MeV (fixed,scale=1000000,gradient)
2020-10-02T09:41:48: Number of temporal par's ..: 1
2020-10-02T09:41:48: Normalization .....: 1 (relative value) (fixed,scale=1,gradient)
2020-10-02T09:41:48: Number of scale par's ....: 0
2020-10-02T09:41:48: === GCTAModelCubeBackground ===
2020-10-02T09:41:48: Name .....: BackgroundModel
2020-10-02T09:41:48: Instruments .....: CTA, HESS, MAGIC, VERITAS
2020-10-02T09:41:48: Observation identifiers ...: all
2020-10-02T09:41:48: Model type .....: "PowerLaw" * "Constant"

```

```
2020-10-02T09:41:48: Number of parameters .....: 4
2020-10-02T09:41:48: Number of spectral par's ..: 3
2020-10-02T09:41:48: Prefactor .....: 0.997217683017508 +/- 0.00234231448712994 [0.01,100] ph/cm2/s/MeV (free,scale=1,gradient)
2020-10-02T09:41:48: Index .....: -0.00232203229406985 +/- 0.00134152892129765 [-5.5] (free,scale=1,gradient)
2020-10-02T09:41:48: PivotEnergy .....: 1000000 MeV (fixed,scale=1000000,gradient)
2020-10-02T09:41:48: Number of temporal par's ..: 1
2020-10-02T09:41:48: Normalization .....: 1 (relative value) (fixed,scale=1,gradient)
2020-10-02T09:41:48:
2020-10-02T09:41:48: +=====
2020-10-02T09:41:48: | Save results |
2020-10-02T09:41:48: +=====
2020-10-02T09:41:48: Model definition file .....: crab_results.xml
2020-10-02T09:41:48: Covariance matrix file ....: NONE
2020-10-02T09:41:48:
2020-10-02T09:41:48: Application "ctlike" terminated after 46 wall clock seconds, consuming 45.6006 seconds of CPU time.
```

#7 - 10/12/2020 11:23 AM - Knödlseder Jürgen

I also checked the speed-up on kepler. Using the vectorised computation, the ctlike run took 111.85 CPU seconds, the former code took 1187.7 CPU seconds, corresponding to a speed-up of a factor of 10.6.

#8 - 10/12/2020 06:15 PM - Knödlseder Jürgen

Using valgrind, I figured out that a substantial amount of time was used for flushing the sparse matrix. The issue was that GModel::eval(GObservation, GMatrixSparse) handled a matrix where the number of columns corresponded to the number of events, and the number of rows to the number of parameters. Therefore, many column operations were required.

Using a transposed matrix, i.e. using a matrix where the number of rows equals to the number of events and the number of columns equals the number of parameters, leads to a considerable reduction of the computing time. Here the result on Mac OS X:

```
2020-10-12T16:12:01: Application "ctlike" terminated after 19 wall clock seconds, consuming 19.4769 seconds of CPU time.
```

which corresponds to a speed-up by a factor of 16. On kepler, the time went down to 55.5 CPU seconds, corresponding to a speed-up factor of 21.

#9 - 10/13/2020 06:39 PM - Knödlseder Jürgen

- % Done changed from 20 to 30

After some additional tweaking of the code, and specifically an attempt of reducing the copying of vectors, and replacing the sparse matrix by a

regular matrix for the response convolution, I got the following results:

System	Former CPU time	New CPU time	Speed-up
Mac OS X	307.8 s	15.4 s	20
Kepler	1187.7 s	53.1 s	22

#10 - 10/14/2020 02:00 PM - Knödlseder Jürgen

I reworked the GIntergrals::polint() method which now operates on vectors. This saved some additional computing time:

System	Former CPU time	New CPU time	Speed-up
Mac OS X	307.8 s	14.6 s	21
Kepler	1187.7 s	51.9 s	23

#11 - 10/15/2020 12:25 AM - Knödlseder Jürgen

Before the branch can be merged into devel the following things need to be done:

- handle all angle exceptions (in particular zeta=0)
- make sure that energy dispersion works
- make sure that unbinned analysis works
- make sure that binned analysis works
- make sure that Chi-squared statistic works
- make sure that all other radial models work

#12 - 10/15/2020 10:01 AM - Knödlseder Jürgen

A scheme needs to be implemented that tracks for which observation spatial model gradients were computed, and for which they are not available. This is at least needed for the transition period, where not all observation types may support the computation of gradients.

#13 - 10/15/2020 03:13 PM - Knödlseder Jürgen

- % Done changed from 30 to 40

I implemented such a scheme and checked that the unbinned analysis works as before.

Here the reference result for the unbinned analysis and a Gaussian model:

```
2020-10-15T12:13:50: +=====+
2020-10-15T12:13:50: | Maximum likelihood optimisation |
2020-10-15T12:13:50: +=====+
2020-10-15T12:28:04: >Iteration 0: -logL=11118401.180, Lambda=1.0e-03
2020-10-15T12:42:16: >Iteration 1: -logL=11118399.292, Lambda=1.0e-03, delta=1.888, step=1.0e+00, max(|grad|)=21.103257 [Sigma:2]
2020-10-15T12:56:32: >Iteration 2: -logL=11118399.291, Lambda=1.0e-04, delta=0.000, step=1.0e+00, max(|grad|)=0.277385 [Sigma:2]
2020-10-15T13:10:35:
2020-10-15T13:10:35: +=====
2020-10-15T13:10:35: | Maximum likelihood optimisation results |
2020-10-15T13:10:35: +=====+
```

```

2020-10-15T13:10:35: === GOptimizerLM ===
2020-10-15T13:10:35: Optimized function value ..: 11118399.291
2020-10-15T13:10:35: Absolute precision .....: 0.005
2020-10-15T13:10:35: Acceptable value decrease ..: 2
2020-10-15T13:10:35: Optimization status .....: converged
2020-10-15T13:10:35: Number of parameters .....: 11
2020-10-15T13:10:35: Number of free parameters ..: 7
2020-10-15T13:10:35: Number of iterations .....: 2
2020-10-15T13:10:35: Lambda .....: 1e-05
2020-10-15T13:10:35: Maximum log likelihood ....: -11118399.291
2020-10-15T13:10:35: Observed events (Nobs) ....: 1927298.000
2020-10-15T13:10:35: Predicted events (Npred) ..: 1927298.000 (Nobs - Npred = 3.99141572415829e-06)
2020-10-15T13:10:35: === GModels ===
2020-10-15T13:10:35: Number of models .....: 2
2020-10-15T13:10:35: Number of parameters .....: 11
2020-10-15T13:10:35: === GModelSky ===
2020-10-15T13:10:35: Name .....: Crab
2020-10-15T13:10:35: Instruments .....: all
2020-10-15T13:10:35: Observation identifiers ...: all
2020-10-15T13:10:35: Model type .....: ExtendedSource
2020-10-15T13:10:35: Model components .....: "RadialGaussian" * "PowerLaw" * "Constant"
2020-10-15T13:10:35: Number of parameters .....: 7
2020-10-15T13:10:35: Number of spatial par's ...: 3
2020-10-15T13:10:35: RA .....: 83.6322949816776 +/- 0.00136695028982184 [-360,360] deg (free,scale=1)
2020-10-15T13:10:35: DEC .....: 22.0133009233324 +/- 0.00126456422869227 [-90,90] deg (free,scale=1)
2020-10-15T13:10:35: Sigma .....: 0.199394483471969 +/- 0.000845557923303607 [0.01,10] deg (free,scale=1)
2020-10-15T13:10:35: Number of spectral par's ...: 3
2020-10-15T13:10:35: Prefactor .....: 5.71182969332522e-16 +/- 3.31633087703927e-18 [1e-23,1e-13] ph/cm2/s/MeV
(free,scale=1e-16,gradient)
2020-10-15T13:10:35: Index .....: -2.47657618741335 +/- 0.00459953482896063 [-5,-0] (free,scale=-1,gradient)
2020-10-15T13:10:35: PivotEnergy .....: 300000 [10000,1000000000] MeV (fixed,scale=1000000,gradient)
2020-10-15T13:10:35: Number of temporal par's ...: 1
2020-10-15T13:10:35: Normalization .....: 1 (relative value) (fixed,scale=1,gradient)
2020-10-15T13:10:35: Number of scale par's ....: 0
2020-10-15T13:10:35: === GCTAModellrfBackground ===
2020-10-15T13:10:35: Name .....: CTABackgroundModel
2020-10-15T13:10:35: Instruments .....: CTA
2020-10-15T13:10:35: Observation identifiers ...: all
2020-10-15T13:10:35: Model type .....: "PowerLaw" * "Constant"
2020-10-15T13:10:35: Number of parameters .....: 4
2020-10-15T13:10:35: Number of spectral par's ...: 3
2020-10-15T13:10:35: Prefactor .....: 0.999337804325637 +/- 0.00114600112209522 [0.001,1000] ph/cm2/s/MeV (free,scale=1,gradient)
2020-10-15T13:10:35: Index .....: -0.000583169536337051 +/- 0.000673206030838423 [-5,5] (free,scale=1,gradient)
2020-10-15T13:10:35: PivotEnergy .....: 1000000 [10000,1000000000] MeV (fixed,scale=1000000,gradient)
2020-10-15T13:10:35: Number of temporal par's ...: 1
2020-10-15T13:10:35: Normalization .....: 1 (relative value) (fixed,scale=1,gradient)
2020-10-15T13:10:35:
2020-10-15T13:10:35: +=====+
2020-10-15T13:10:35: | Save results |
2020-10-15T13:10:35: +=====+
2020-10-15T13:10:35: Model definition file ....: crab_results_unbinned_ref.xml
2020-10-15T13:10:35: Covariance matrix file ....: NONE
2020-10-15T13:10:35:
2020-10-15T13:10:35: Application "ctlike" terminated after 3405 wall clock seconds, consuming 3387.85 seconds of CPU time.

```

And here the result with the current code:

```

2020-10-15T09:59:03: +=====
2020-10-15T09:59:03: | Maximum likelihood optimisation |
2020-10-15T09:59:03: +=====
2020-10-15T10:13:02: >Iteration 0: -logL=11118401.180, Lambda=1.0e-03
2020-10-15T10:27:00: >Iteration 1: -logL=11118399.292, Lambda=1.0e-03, delta=1.888, step=1.0e+00, max(|grad|)=21.103227 [Sigma:2]
2020-10-15T10:40:53: >Iteration 2: -logL=11118399.291, Lambda=1.0e-04, delta=0.000, step=1.0e+00, max(|grad|)=0.277384 [Sigma:2]
2020-10-15T10:54:25:
2020-10-15T10:54:25: +=====
2020-10-15T10:54:25: | Maximum likelihood optimisation results |
2020-10-15T10:54:25: +=====
2020-10-15T10:54:25: === GOptimizerLM ===
2020-10-15T10:54:25: Optimized function value ..: 11118399.291
2020-10-15T10:54:25: Absolute precision .....: 0.005
2020-10-15T10:54:25: Acceptable value decrease ..: 2
2020-10-15T10:54:25: Optimization status .....: converged

```

```

2020-10-15T10:54:25: Number of parameters .....: 11
2020-10-15T10:54:25: Number of free parameters .: 7
2020-10-15T10:54:25: Number of iterations .....: 2
2020-10-15T10:54:25: Lambda .....: 1e-05
2020-10-15T10:54:25: Maximum log likelihood ....: -11118399.291
2020-10-15T10:54:25: Observed events (Nobs) ...: 1927298.000
2020-10-15T10:54:25: Predicted events (Npred) ..: 1927298.000 (Nobs - Npred = 2.59648077189922e-05)
2020-10-15T10:54:25: === GModels ===
2020-10-15T10:54:25: Number of models .....: 2
2020-10-15T10:54:25: Number of parameters .....: 11
2020-10-15T10:54:25: === GModelSky ===
2020-10-15T10:54:25: Name .....: Crab
2020-10-15T10:54:25: Instruments .....: all
2020-10-15T10:54:25: Observation identifiers ...: all
2020-10-15T10:54:25: Model type .....: ExtendedSource
2020-10-15T10:54:25: Model components .....: "RadialGaussian" * "PowerLaw" * "Constant"
2020-10-15T10:54:25: Number of parameters .....: 7
2020-10-15T10:54:25: Number of spatial par's ...: 3
2020-10-15T10:54:25: RA .....: 83.6322949816869 +/- 0.00136695028999332 [-360,360] deg (free,scale=1,gradient)
2020-10-15T10:54:25: DEC .....: 22.0133009240316 +/- 0.00126456422883072 [-90,90] deg (free,scale=1,gradient)
2020-10-15T10:54:25: Sigma .....: 0.199394483861209 +/- 0.000845557920483975 [0.01,10] deg (free,scale=1,gradient)
2020-10-15T10:54:25: Number of spectral par's ...: 3
2020-10-15T10:54:25: Prefactor .....: 5.7118296785183e-16 +/- 3.3163308793508e-18 [1e-23,1e-13] ph/cm2/s/MeV
(free,scale=1e-16,gradient)
2020-10-15T10:54:25: Index .....: -2.47657617655134 +/- 0.00459953479184992 [-5,-0] (free,scale=-1,gradient)
2020-10-15T10:54:25: PivotEnergy .....: 300000 [10000,1000000000] MeV (fixed,scale=1000000,gradient)
2020-10-15T10:54:25: Number of temporal par's ...: 1
2020-10-15T10:54:25: Normalization .....: 1 (relative value) (fixed,scale=1,gradient)
2020-10-15T10:54:25: Number of scale par's ...: 0
2020-10-15T10:54:25: === GCTAModellIrfBackground ===
2020-10-15T10:54:25: Name .....: CTABackgroundModel
2020-10-15T10:54:25: Instruments .....: CTA
2020-10-15T10:54:25: Observation identifiers ...: all
2020-10-15T10:54:25: Model type .....: "PowerLaw" * "Constant"
2020-10-15T10:54:25: Number of parameters .....: 4
2020-10-15T10:54:25: Number of spectral par's ...: 3
2020-10-15T10:54:25: Prefactor .....: 0.999337771992533 +/- 0.00114600107737938 [0.001,1000] ph/cm2/s/MeV (free,scale=1,gradient)
2020-10-15T10:54:25: Index .....: -0.000583194436610381 +/- 0.00067320600775784 [-5,5] (free,scale=1,gradient)
2020-10-15T10:54:25: PivotEnergy .....: 1000000 [10000,1000000000] MeV (fixed,scale=1000000,gradient)
2020-10-15T10:54:25: Number of temporal par's ...: 1
2020-10-15T10:54:25: Normalization .....: 1 (relative value) (fixed,scale=1,gradient)
2020-10-15T10:54:25:
2020-10-15T10:54:25: +=====
2020-10-15T10:54:25: | Save results |
2020-10-15T10:54:25: +=====
2020-10-15T10:54:25: Model definition file .....: crab_results.xml
2020-10-15T10:54:25: Covariance matrix file ....: NONE
2020-10-15T10:54:25:
2020-10-15T10:54:25: Application "ctlike" terminated after 3323 wall clock seconds, consuming 3313.09 seconds of CPU time.

```

Results and computing time are identical.

Here a comparison of various different spatial models, including also an unbinned analysis of the Gaussian model. Quantities with `_r` are reference results for the old code, quantities with `_n` are results with the modified code.

Method/Model	CPU_r	Iter_r	logL_r	CPU_n	Iter_n	logL_n	Comments
Unbinned/Gaussian	3387.85	2	-11118399.291	3313.09	2	-11118399.291	identical
Binned/Gaussian	2054.96	2	118107.706	1974.85	2	118107.706	identical
Stacked/Chisqr/Gaussian	2.49	3	-2006.024	2.47	3	-2006.024	identical (heavily rebinned)
Stacked/Disk	35.21	2	156240.662	18.46	3	156240.518	result differs
Stacked/Gaussian	307.80	2	118106.615	15.51	2	118106.615	optimised code
Stacked/Ring	689.32	31	125762.471	158.75	18	125753.377	result differs
Stacked/Shell	143.82	4	127541.952	431.63	14	127539.10	result differs
Stacked/Profile	404.67	2	118004.032	84.99	2	118004.030	result differs slightly
Stacked/Elliptical disk	1187.04	3	83893.643	1128.28	3	83893.643	identical
Stacked/Elliptical Gaussian	3098.16	3	67833.148	3188.62	3	67833.148	identical
Stacked/Diffuse map	90.26	3	56782.544	441.67	3	56782.544	identical, no caching in vector code

One difference with the new vectorised code for radial models is that there is no transition point added in the integration step. I therefore switched on the transition point to see whether it makes a difference. Results are given below.

Model	CPU_r	Iter_r	logL_r	CPU_n	Iter_n	logL_n	Speed-up	Comments
Disk	35.21	2	156240.662	16.09	2	156240.662	2.2	identical results
Gaussian	307.80	2	118106.615	18.07	2	118106.615	17.0	identical results
Ring	689.32	31	125762.471	296.11	31	125762.471	2.3	identical results
Shell	143.82	4	127541.952	58.56	4	127541.952	2.5	identical results
Profile	404.67	2	118004.032	113.81	2	118004.032	3.6	identical results

With the transition point the results are identical, and the fitting is more stable for some of the models, hence it's better to leave it on. However the fitting takes a bit longer, and the speed-up for the Gaussian model, which was already exact without a transition point, is smaller.

It should be checked whether a transition point is actually needed for the other models once analytical gradients are implemented.

- % Done changed from 40 to 50

I started to check the energy dispersion. Here the result of a reference run:

```
2020-10-16T09:54:44: +=====
2020-10-16T09:54:44: | Maximum likelihood optimisation |
2020-10-16T09:54:44: +=====
2020-10-16T10:42:36: >Iteration 0: -logL=-122525.568, Lambda=1.0e-03
2020-10-16T11:30:31: >Iteration 1: -logL=-122531.428, Lambda=1.0e-03, delta=5.860, step=1.0e+00, max(|grad|)=27.397503 [Sigma:2]
2020-10-16T12:17:50: >Iteration 2: -logL=-122531.429, Lambda=1.0e-04, delta=0.000, step=1.0e+00, max(|grad|)=0.171484 [Sigma:2]
2020-10-16T13:05:54:
2020-10-16T13:05:54: +=====
2020-10-16T13:05:54: | Maximum likelihood optimisation results |
2020-10-16T13:05:54: +=====
2020-10-16T13:05:54: === GOptimizerLM ===
2020-10-16T13:05:54: Optimized function value ...: -122531.429
2020-10-16T13:05:54: Absolute precision .....: 0.005
2020-10-16T13:05:54: Acceptable value decrease ..: 2
2020-10-16T13:05:54: Optimization status .....: converged
2020-10-16T13:05:54: Number of parameters .....: 11
2020-10-16T13:05:54: Number of free parameters ..: 7
2020-10-16T13:05:54: Number of iterations .....: 2
2020-10-16T13:05:54: Lambda .....: 1e-05
2020-10-16T13:05:54: Maximum log likelihood .....: 122531.429
2020-10-16T13:05:54: Observed events (Nobs) ...: 785964.000
2020-10-16T13:05:54: Predicted events (Npred) ...: 785964.000 (Nobs - Npred = 0.00023000470764935)
2020-10-16T13:05:54: === GModels ===
2020-10-16T13:05:54: Number of models .....: 2
2020-10-16T13:05:54: Number of parameters .....: 11
2020-10-16T13:05:54: === GModelSky ===
2020-10-16T13:05:54: Name .....: Crab
2020-10-16T13:05:54: Instruments .....: all
2020-10-16T13:05:54: Observation identifiers ...: all
2020-10-16T13:05:54: Model type .....: ExtendedSource
2020-10-16T13:05:54: Model components .....: "RadialGaussian" * "PowerLaw" * "Constant"
2020-10-16T13:05:54: Number of parameters .....: 7
2020-10-16T13:05:54: Number of spatial par's ...: 3
2020-10-16T13:05:54: RA .....: 83.6333193130751 +/- 0.00135238184050406 [-360,360] deg (free,scale=1)
2020-10-16T13:05:54: DEC .....: 22.0154036092787 +/- 0.00124858809704022 [-90,90] deg (free,scale=1)
2020-10-16T13:05:54: Sigma .....: 0.199957055498399 +/- 0.000845114002169637 [0.01,10] deg (free,scale=1)
2020-10-16T13:05:54: Number of spectral par's ...: 3
2020-10-16T13:05:54: Prefactor .....: 5.63575354021804e-16 +/- 3.21590631991281e-18 [1e-23,1e-13] ph/cm2/s/MeV
(free,scale=1e-16,gradient)
2020-10-16T13:05:54: Index .....: -2.48593583530845 +/- 0.00453433707882368 [-5,-0] (free,scale=-1,gradient)
2020-10-16T13:05:54: PivotEnergy .....: 300000 [10000,1000000000] MeV (fixed,scale=1000000,gradient)
2020-10-16T13:05:54: Number of temporal par's ...: 1
2020-10-16T13:05:54: Normalization .....: 1 (relative value) (fixed,scale=1,gradient)
2020-10-16T13:05:54: Number of scale par's ....: 0
2020-10-16T13:05:54: === GCTAModelCubeBackground ===
2020-10-16T13:05:54: Name .....: BackgroundModel
2020-10-16T13:05:54: Instruments .....: CTA, HESS, MAGIC, VERITAS
2020-10-16T13:05:54: Observation identifiers ...: all
2020-10-16T13:05:54: Model type .....: "PowerLaw" * "Constant"
2020-10-16T13:05:54: Number of parameters .....: 4
2020-10-16T13:05:54: Number of spectral par's ...: 3
2020-10-16T13:05:54: Prefactor .....: 0.99963165325343 +/- 0.00234123613800205 [0.01,100] ph/cm2/s/MeV (free,scale=1,gradient)
2020-10-16T13:05:54: Index .....: -0.00212242798329811 +/- 0.0013415027282217 [-5,5] (free,scale=1,gradient)
2020-10-16T13:05:54: PivotEnergy .....: 1000000 MeV (fixed,scale=1000000,gradient)
2020-10-16T13:05:54: Number of temporal par's ...: 1
2020-10-16T13:05:54: Normalization .....: 1 (relative value) (fixed,scale=1,gradient)
2020-10-16T13:05:54:
2020-10-16T13:05:54: +=====
2020-10-16T13:05:54: | Save results |
2020-10-16T13:05:54: +=====
2020-10-16T13:05:54: Model definition file .....: crab_results_ref.xml
2020-10-16T13:05:54: Covariance matrix file ....: NONE
2020-10-16T13:05:54:
2020-10-16T13:05:54: Application "ctlike" terminated after 11471 wall clock seconds, consuming 11385.7 seconds of CPU time.
```

After fixing a small bug in GResponse with the indices for the energy dispersion computation I got the following result using the new code. The results are identical. This completes all the tests that were enumerated above.

```

2020-10-16T18:25:59: +=====
2020-10-16T18:25:59: | Maximum likelihood optimisation |
2020-10-16T18:25:59: +=====
2020-10-16T19:10:23: >Iteration 0: -logL=-122525.568, Lambda=1.0e-03
2020-10-16T19:56:37: >Iteration 1: -logL=-122531.428, Lambda=1.0e-03, delta=5.860, step=1.0e+00, max(|grad|)=27.397503 [Sigma:2]
2020-10-16T20:42:55: >Iteration 2: -logL=-122531.429, Lambda=1.0e-04, delta=0.000, step=1.0e+00, max(|grad|)=0.171484 [Sigma:2]
2020-10-16T21:29:38:
2020-10-16T21:29:38: +=====
2020-10-16T21:29:38: | Maximum likelihood optimisation results |
2020-10-16T21:29:38: +=====
2020-10-16T21:29:38: === GOptimizerLM ===
2020-10-16T21:29:38: Optimized function value ...: -122531.429
2020-10-16T21:29:38: Absolute precision .....: 0.005
2020-10-16T21:29:38: Acceptable value decrease ..: 2
2020-10-16T21:29:38: Optimization status .....: converged
2020-10-16T21:29:38: Number of parameters .....: 11
2020-10-16T21:29:38: Number of free parameters ..: 7
2020-10-16T21:29:38: Number of iterations .....: 2
2020-10-16T21:29:38: Lambda .....: 1e-05
2020-10-16T21:29:38: Maximum log likelihood ....: 122531.429
2020-10-16T21:29:38: Observed events (Nobs) ...: 785964.000
2020-10-16T21:29:38: Predicted events (Npred) ...: 785964.000 (Nobs - Npred = 0.00023000470764935)
2020-10-16T21:29:38: === GModels ===
2020-10-16T21:29:38: Number of models .....: 2
2020-10-16T21:29:38: Number of parameters .....: 11
2020-10-16T21:29:38: === GModelSky ===
2020-10-16T21:29:38: Name .....: Crab
2020-10-16T21:29:38: Instruments .....: all
2020-10-16T21:29:38: Observation identifiers ...: all
2020-10-16T21:29:38: Model type .....: ExtendedSource
2020-10-16T21:29:38: Model components .....: "RadialGaussian" * "PowerLaw" * "Constant"
2020-10-16T21:29:38: Number of parameters .....: 7
2020-10-16T21:29:38: Number of spatial par's ...: 3
2020-10-16T21:29:38: RA .....: 83.6333193130751 +/- 0.00135238184050406 [-360,360] deg (free,scale=1,gradient)
2020-10-16T21:29:38: DEC .....: 22.0154036092787 +/- 0.00124858809704022 [-90,90] deg (free,scale=1,gradient)
2020-10-16T21:29:38: Sigma .....: 0.199957055498399 +/- 0.000845114002169637 [0.01,10] deg (free,scale=1,gradient)
2020-10-16T21:29:38: Number of spectral par's ...: 3
2020-10-16T21:29:38: Prefactor .....: 5.63575354021804e-16 +/- 3.21590631991281e-18 [1e-23,1e-13] ph/cm2/s/MeV
(free,scale=1e-16,gradient)
2020-10-16T21:29:38: Index .....: -2.48593583530845 +/- 0.00453433707882368 [-5,-0] (free,scale=-1,gradient)
2020-10-16T21:29:38: PivotEnergy .....: 300000 [10000,1000000000] MeV (fixed,scale=1000000,gradient)
2020-10-16T21:29:38: Number of temporal par's ...: 1
2020-10-16T21:29:38: Normalization .....: 1 (relative value) (fixed,scale=1,gradient)
2020-10-16T21:29:38: Number of scale par's ....: 0
2020-10-16T21:29:38: === GCTAModelCubeBackground ===
2020-10-16T21:29:38: Name .....: BackgroundModel
2020-10-16T21:29:38: Instruments .....: CTA, HESS, MAGIC, VERITAS
2020-10-16T21:29:38: Observation identifiers ...: all
2020-10-16T21:29:38: Model type .....: "PowerLaw" * "Constant"
2020-10-16T21:29:38: Number of parameters .....: 4
2020-10-16T21:29:38: Number of spectral par's ...: 3
2020-10-16T21:29:38: Prefactor .....: 0.996936165325343 +/- 0.00234123613800205 [0.01,100] ph/cm2/s/MeV (free,scale=1,gradient)
2020-10-16T21:29:38: Index .....: -0.00212242798329811 +/- 0.0013415027282217 [-5,5] (free,scale=1,gradient)
2020-10-16T21:29:38: PivotEnergy .....: 1000000 MeV (fixed,scale=1000000,gradient)
2020-10-16T21:29:38: Number of temporal par's ...: 1
2020-10-16T21:29:38: Normalization .....: 1 (relative value) (fixed,scale=1,gradient)
2020-10-16T21:29:38:
2020-10-16T21:29:38: +=====
2020-10-16T21:29:38: | Save results |
2020-10-16T21:29:38: +=====
2020-10-16T21:29:38: Model definition file ....: crab_results_new.xml
2020-10-16T21:29:38: Covariance matrix file ....: NONE
2020-10-16T21:29:38:
2020-10-16T21:29:38: Application "ctlike" terminated after 11019 wall clock seconds, consuming 10871.2 seconds of CPU time.

```

#17 - 10/16/2020 11:37 PM - Knödlseder Jürgen

I also implemented the vectorised code for the unbinned and gaussian likelihoods in GObservation. Below the updated comparison for these two cases. Everything seems to work as expected.

Method/Model	CPU_r	Iter_r	logL_r	CPU_n	Iter_n	logL_n	Comments
Unbinned/Gaussian	3387.85	2	-11118399.291	3271.67	2	-11118399.291	identical
Stacked/Chisqr/Gaussian	2.49	3	-2006.024	0.20	3	-2006.024	identical

#18 - 10/17/2020 12:50 AM - Knödlseder Jürgen

- % Done changed from 50 to 60

I thought I had finished with the code to be merged in, but the LAT unit test is taking a very long time. Apparently the vectorised response computation has some kind of side effect on the LAT response computation.

#19 - 10/17/2020 02:58 PM - Knödlseder Jürgen

- Status changed from In Progress to Pull request
- % Done changed from 60 to 90

The LAT interface had not virtual methods GLATResponse::irf_ptsrc and GLATResponse::irf_diffuse that are needed for the vectorized response to work correctly. I added such methods and the unit tests were successful.

#20 - 10/18/2020 12:45 AM - Knödlseder Jürgen

- Status changed from Pull request to In Progress

The code integration revealed a vulnerability with stack of model pointers m_spat_pars_with_gradients that is currently used to check whether a parameter has an analytical gradient. It can happen that a model gets out of scope, and hence the pointer becomes invalid. It seems safer to store a string, composed of the model name and the parameter name.

#21 - 10/18/2020 10:51 AM - Knödlseder Jürgen

- Status changed from In Progress to Closed
- % Done changed from 90 to 100

I fixed the problem and merged the code into devel.

#22 - 10/19/2020 02:40 PM - Knödlseder Jürgen

I finally did another change that relates to this issue and that consists of using the scalar Phi integrator in case that no gradients are requested. This led to the following speed ups:

Model	CPU_r	Iter_r	logL_r	CPU_n	Iter_n	logL_n	Speed-up	Comments
Disk	35.21	2	156240.662	14.81	2	156240.662	2.4	identical results
Gaussian	307.80	2	118106.615	16.55	2	118106.615	18.6	identical results

Ring	689.32	31	125762.471	257.40	31	125762.471	2.7	identical results
Shell	143.82	4	127541.952	48.47	4	127541.952	3.0	identical results
Profile	404.67	2	118004.032	97.40	2	118004.032	4.2	identical results

Files

gammalib_tn0003_response.pdf

163 KB

04/13/2020

Knödlseder Jürgen