Investigations of the scalability of the H.E.S.S. software for CTA Monte Carlo studies

Dirk Lennarz[1]
III. Physikalisches Institut B, RWTH Aachen University, 52056 Aachen, Germany
dirk.lennarz@physik.rwth-aachen.de

Abstract

The SASH software package was investigated with respect to its performance as a function of the number of telescopes in a planned array of CTA telescopes. Looking at Monte Carlo simulations of four different telescope setups the time to read files and access the events was measured. It could be shown that including more events into a file decreases the time to read a certain amount of events. For one telescope setup the time to read or access the events is linear to the number of files or accessed events. Comparing the four setups it could be shown that the time to read a file depends on the number of telescopes, but that the time to access events scales linearly with it. Presumably more detailed investigations are needed before SASH can become the software for the future array of CTA telescopes.

1. Introduction to gamma-ray astronomy

In recent years very high energy (VHE) γ-ray astronomy has opened a complete new window to the universe of the highest energies. Besides spectacular new insights like the first image of the Milky Way in VHE γ-rays also surprising results came up for example the discovery of "dark sources" which emit VHE γ-rays but are not visible in other wavelength regimes.

1.1. Cherenkov telescopes

The recent breakthroughs were achieved by Cherenkov telescopes like H.E.S.S., MAGIC and VERITAS in the northern hemisphere and CANGAROO in the southern hemisphere. A VHE γ entering the atmosphere will at some point interact with an atmospheric nuclei and generate an electromagnetic shower. As the electrons and positrons travel faster than the speed of light in air they will emit Cherenkov light. The light will typically illuminate an area of about 250 m diameter on the ground and can be used to reconstruct the celestial origin of the incoming particle. All current instruments use multiple telescopes to view at the shower from different angels and use finely-pixellated images. These "stereoscopic" telescopes have improved sensitivity for VHE γ-rays by more than two orders of magnitude since the first VHE γ-ray image of the Crab Nebula in 1989.

1.2. The Cherenkov Telescope Array

The flux of VHE γ-rays gets significantly smaller at higher energies so a large effective detection area is needed in order to boost the event rates. The Cherenkov Telescope Array (CTA) is intended for becoming the next generation VHE γ-ray telescope infrastructure, aiming at providing a so far unseen sensitivity. It would succeed the very successful European experiments H.E.S.S. and MAGIC. Demonstrated technologies and the combined know-how from the convergence of two experiments could be used to create a new facility on an unprecedented scale.

At the moment CTA has entered into a five year lasting R&D and prototype phase. It is planned that CTA enters in construction phase in 2012. A letter of intent for CTA is still underway. The optimal design for the CTA array is currently investigated. It will most likely consist of a mix of

[1]Summer Student 2008, Deutsche Elektronen-Synchrotron DESY, Zeuthen

Table 1

Averaged values over all files

	HESS I	HESS II	CTA 41	CTA 97
kB/event	0.94	0.70	4.35	10.3
MB/file	5.57	8.66	89.0	200
events/file	6051.5	12739.3	20935.7	20000

different telescopes types. One of the possible designs is to have a central part with large telescopes, surrounded by a ring of medium sized and lastly surrounded by a third ring of smaller telescopes.

2. The H.E.S.S. software framework

A new software and data format will be needed for CTA. There exists a large variety of telescope simulations and data formats from different experiments. The analysis described in this report deals with the performance of SASH, the Storage and Analysis Software at H.E.S.S., with respect to Monte Carlo samples simulating the behaviour of a large number of telescopes planned for CTA.

2.1. The SASH software package

SASH provides classes for data storage and analysis in the H.E.S.S. software framework. It is written in C++, but heavily based on the `ROOT` framework from CERN. For the upcoming analysis the version berlin-0.7.13 of the SASH software was used, together with the `ROOT` version 5.16.00 and version 4.1.2 of the gcc compiler.

2.2. Data simulation and storage

Shower simulations in the H.E.S.S. experiment are done for example via the CORSIKA software package. Once the shower is simulated the Cherenkov photons are calculated. In the next step the telescope array is simulated and only the part of the light cone that actually reaches a telescope is kept. At this level the data is considered to be raw data.

For faster and simpler high level analysis the information is further reduced to so called DST level. The data undergoes a cleaning which consists in the elimination of isolated pixels and of pixel intensities below a threshold. The Hillas parameters, used for the reconstruction of the incoming particle, are calculated.

The information from raw data and DST files is stored in `ROOT` files using TTrees.

3. Used simulation data

For this work Monte Carlo simulations on DST level[1] for four different telescope setups with 4, 5, 41 and 97 telescopes were used. The 4 telescopes corresponds to the current H.E.S.S. setup (referred to as Phase I or HESS I), while the setup with 5 telescopes includes a planned update of the current H.E.S.S. setup (referred to as Phase II or HESS II) with one large telescope in the middle of the four others. For the time measurement the actual design of the two used CTA setups is of minor importance because only the number of telescopes matters. The MC input was a γ-ray point source with a E^{-2} spectrum.

As it can be seen in Fig. 1 the file size is correlated to the number of events in the file. For HESS I the file size changes over two orders of magnitude, going down as low as 400 kB. For the other telescope setups the file size is more homogenous. It is not known why the simulation for HESS I is so different from the other three telescope setups. For the HESS I files another oddity was observed: there is a clear correlation between the run number and the file size (see Fig. 2).

From this data the average size of one event, file size and the events per files was calculated by dividing the total number of bytes, events and files respectively (see Tab. 1). It can be seen that the event size decreases between HESS I and HESS II. This is somehow unexpected because one telescope more can contribute data to the event.

There can be various explanations for this behaviour. The HESS I files contain only one run per file, while for the other three setups two runs

[1] By K. Bernloehr from Heidelberg, DST conversion done with Berlin software

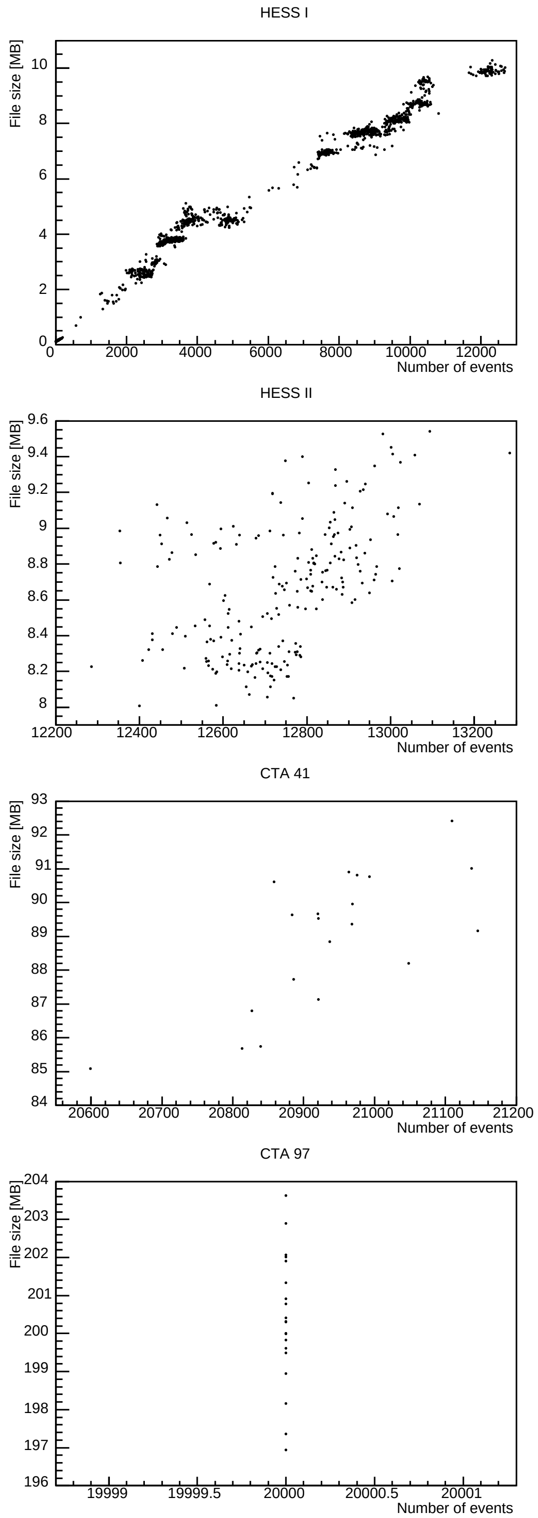

Fig. 1.— Correlation between the file size and the number of events in the file.

were included in one file. This could reduce the file-overhead, because information about for example the generating Monte Carlo has only to be stored once. On the other hand this could be just an effect of the compression `ROOT` does when storing TTrees. Another explanation is provided by the studies done in the last paragraph. It will be shown, that the medium number of telescopes with data changed from 2.8 for HESS I to 1.8 for HESS II, which will be explained in more detail in that paragraph. This could lead to less data that has to be saved for every event.

The increasing average file size from CTA 41 to 97 is however quiet puzzling, because the number of telescopes with data is almost equal (see last paragraph). The limited range of the Cherenkov light cone ($\approx$ 250m) causes the physical information of one air shower to be spread over only a few telescopes (distance 100m). Hence, after a certain size of the array the event size on DST level should stay more or less constant. One possible explanation could be that also other parameters, like the size of the telescopes and with that the medium number of pixels per telescope where varied between the two setups.

4. Time for reading files

In this paragraph "reading" of files means that the events in the files are made available for analysis usage. In the language of `ROOT` one would say that the events from the trees were chained, but are not accessed yet. A DST and a

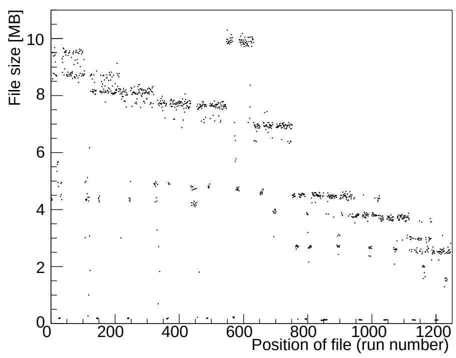

Fig. 2.— For HESS I the file size is correlated to the run number assigned to the file.

run Sash::DataSet were created and a dependency added between them. Files are then added to the data set.

The files were read in 10% portions of the total number of available files and the necessary CPU time was measured. That means that first 10% of the files were read, afterwards the data set was deleted, then 20% of the data were read and so on. All this was repeated 10 times in order to check on the stability of the timing. These measurements were then averaged and the RMS taken as error. The final ten data points were finally fitted with a pol1 function from `ROOT`. The fits are shown in Fig. 4. The χ^2 per NDF is in general to small so the errors are overestimated. It can be argued that the RMS is not the correct error, but the fits are nevertheless meaningful.

Comparing the HESS I plot in Fig. 4 with the correlation between run number and file size from Fig. 2 it becomes obvious, that the reading of the file does not depend on the file size. Otherwise the first files should be read slower than last ones. This file size independence is somehow logical, because the events are so far only "chained". The time for this should not depend on the number of events in the file.

Keeping in mind that the file size does not influence the time to read a file it is really astonishing that the time to read one file does depend on the number of telescopes. This is shown in Fig. 3. The scaling effect is not expected and as the reason for this remains unknown, further investigation is needed to clarify what causes this delay

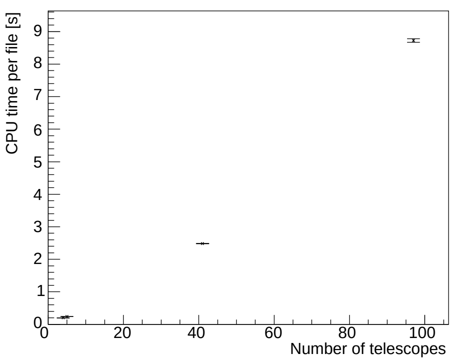

Fig. 3.— CPU time for reading one files as a function of number of telescopes

for larger telescope arrays.

The CPU time was also measured as a function of the read events (see Fig. 5). With the values found in Tab. 1 the before fitted slope was converted via the medium number of events per file. This worked well for HESS II, CTA 41 and CTA 97. At HESS I we can see an effect from the fact shown in Fig. 2: the first read files are large so the number of events increases faster than the expectation from the medium number of events per file, while smaller files are read later.

Comparing HESS I to HESS II it becomes clear, that putting more events into one file reduces the time necessary to read a million events. Due to different event sizes this hold true even if only the first 500 files from HESS I, which have a comparable file size to the HESS II files, are compared.

The relative values comparing the different telescope setups are very stable against changes. It was found that the absolut value changes slightly with the used CPU. Sometimes a dramatic change up to 50% was observed, presumingly due to heavy load on the file system. Because of the great influence on the time these measurements could be easily sorted out.

5. Time for looping through events

In the next step the before only chained events were accessed. In the H.E.S.S. framework the analysis is written in a so called Sash::Maker which is compiled to a library and was then used in CINT. The maker consists in just getting the Sash::HESSArray and reading from it the Sash::McTrueShower. The true energy from the Monte Carlo was extracted from it, without using it otherwise. This is a value which is stored one time per event.

Using the compiled maker the loop was then done with a compiled program in CINT. Data sets were created and files with a sufficient number of events were read in order to reach the desired amount of maximal number of events. After that a data iterator and a chain were created and using the maker all the events were processed. This was done for all four telescope setups. Only the time of the event processing was measured. Because of the huge amount of time it took for CTA 41 the timing was only repeated 5 times. The averaging and fitting was done the same way as in the last

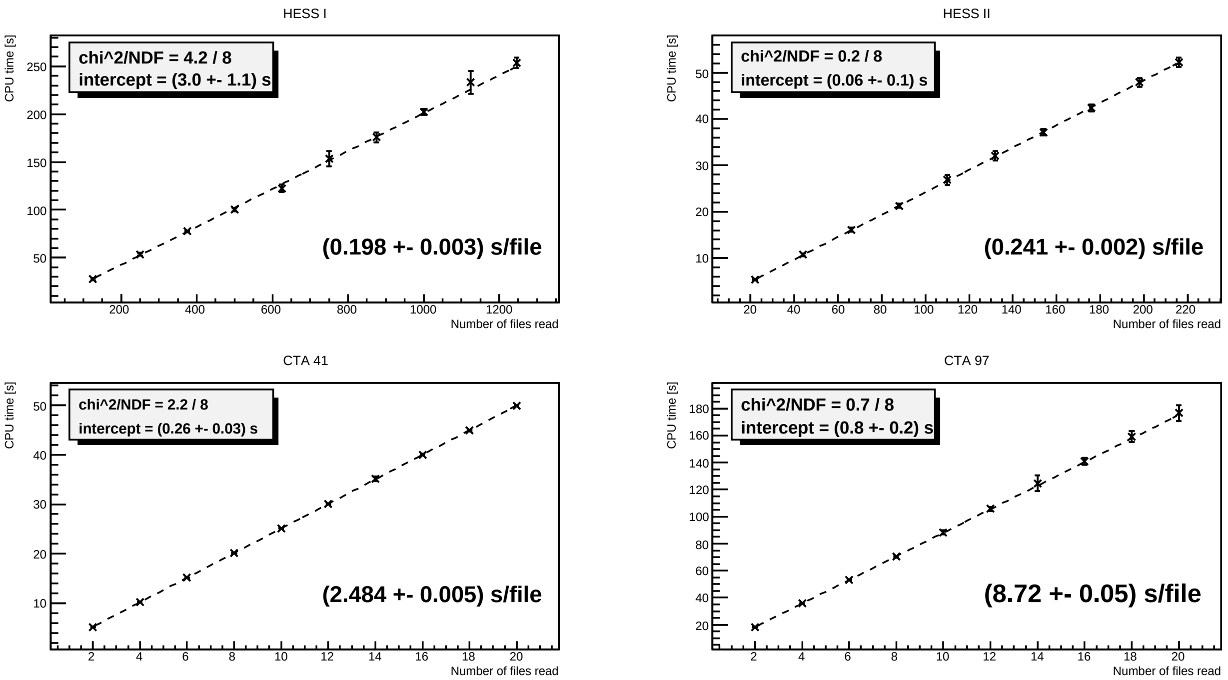

Fig. 4.— CPU time for reading files as a function of the number of files

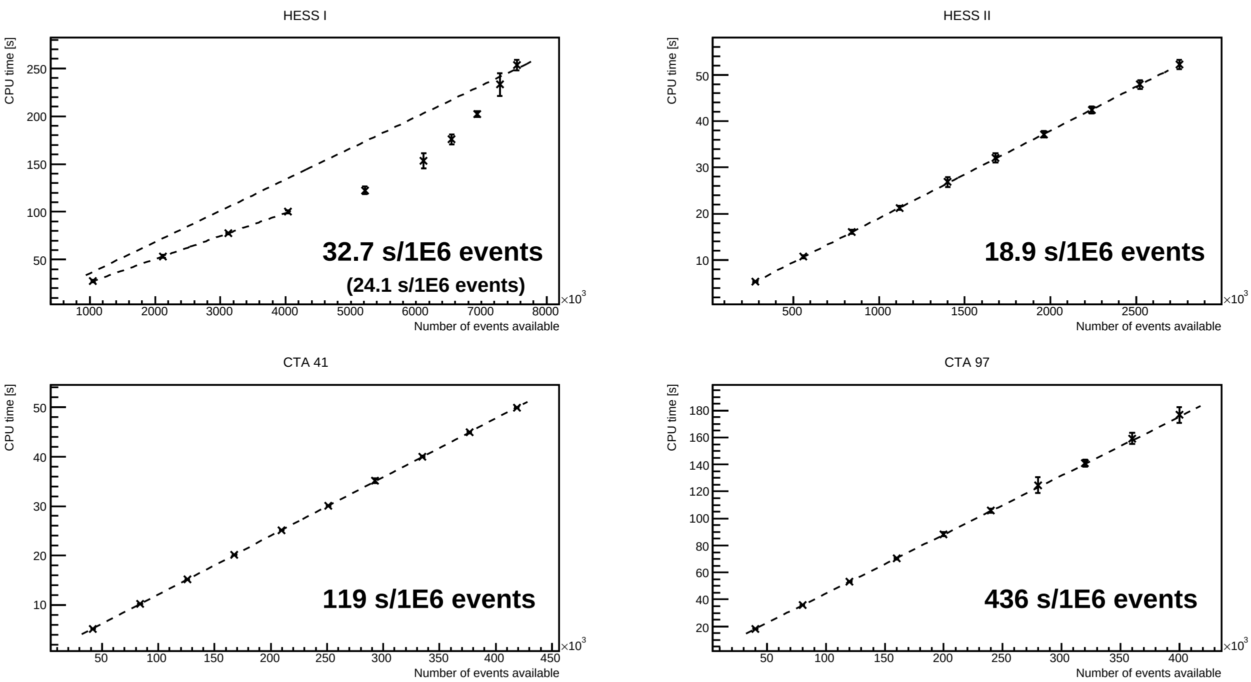

Fig. 5.— CPU time for reading files as a function of the number of events. The dashed lines are the fit from the time to open a file converted by the medium number of events per file. For HESS I the first 4 points were fitted separately.

paragraph. The fits for each telescope setup are shown in the appendix (see Fig. 8). The χ^2 per NDF is quiet acceptable in most cases, despite of the known issue with the RMS error.

In Fig. 6 it can be seen that the CPU time scales linearly with the number of telescopes. This is evident even if the fit is not perfect. The negative interception with the y-axis could be an indicator to a minimal time that is constant for a small number of telescopes and that the linear increase starts at some later point.

This scaling behaviour is again a rather surprising outcome. First of all it indicates that, although only one value which is stored once per event but not once per telescope is accessed, the information of the event is always read completely. The linear scaling could be due to the different event sizes. What was not taken into account here is the fact that to loop the same number of events a different number of files was looped through for all four setups.

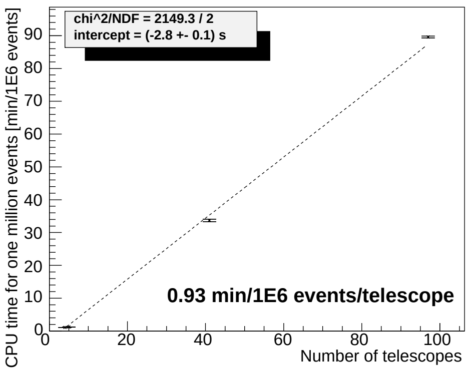

Fig. 6.— CPU time for looping through one million events as a function of number of telescopes

6. Number of triggered telescopes

The maker from the last paragraph was modified so that it figures out how many telescopes are in the setup in order to know if it is a HESS I, HESS II, CTA 41 or CTA 97 event. Then the Sash::EventHeader was loaded and the number of triggered telescopes and of telescopes with data accessed. All the information was filed into a ROOT tree.

Comparing the number of triggered telescopes

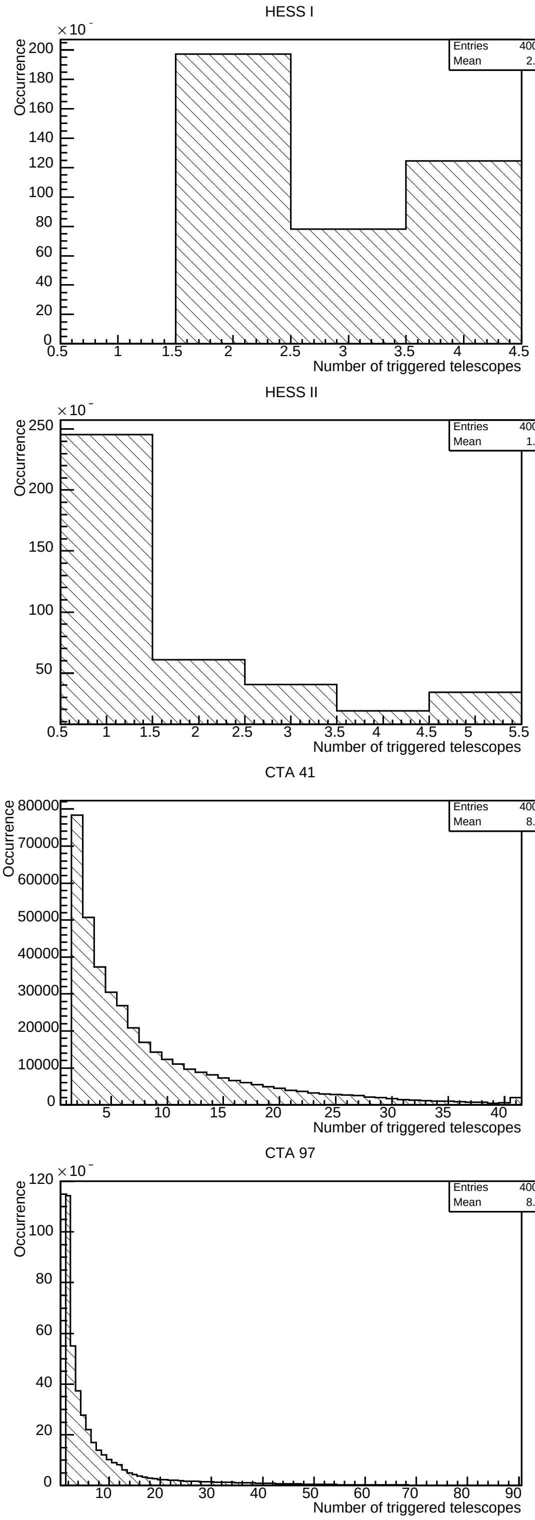

Fig. 7.— Distribution of triggered telescopes

to the number of telescopes with data it was found that they are most of the time identically, though there are 17 events in HESS I where this is not the case.

In Fig. 7 the number of the triggered telescopes was written to a histogram for 400,000 events for each telescope setup. The medium number has not changed from CTA 41 to CTA 97 which is, due to the limited area of the cherenkov light cone, what one would expect.

The data from HESS I was only taken if two telescopes triggered in order to be able to actually reconstruct the direction (stereoscopic). In HESS II, which should be updated with a bigger telescope in the middle, the big telescope was most likely allowed to trigger by itself. For the other two CTA setups there are again no events triggered by only one telescope. There seems to be an excess of CTA 41 events where all 41 telescopes triggered. Maybe even more surprisingly is the fact that there is a CTA 97 event where 90 telescopes have data.

7. Conclusions

As illustrated in this report the SASH software needs more detailed investigation before it can become the software for a future array of CTA telescopes. It could be possible that some changes on the code become necessary.

It was clearly seen that putting more events into one file reduces the time to read a certain amount of events. Maybe this even leads to a reduced size of one event on the hard disk. Why the reading of files scales with the number of telescopes still remains mysterious and should be checked.

The time to loop through events depends linearly on the number of telescopes. As the complete event seems to be read to memory, the different sizes of one event will surely have an influence on that. Also on this point further investigation is needed.

Acknowledgments

The author likes to thank Christian Spiering for the possibility to work in the Zeuthen astroparticle group. A big thank goes to my supervisor Stefan Schlenstedt for his excellent support and the many fruitful discussions during the project and to Heike Prokoph for her help to understand the H.E.S.S. software. I also want to thank Ullrich Schwanke for creating the topic of the project and the discussions about the results.

Finally I would like to thank all the people who made this Summer Student program possible in the first place, there to mention especially Sabine Baer, Silvia Rueckert, Thomas Naumann, Andrea Schrader and of course Karl Jansen for organising this wonderful trip to Hamburg. And last but not least I would like to thank all the other Summer Students for making me pass such a great time in Zeuthen!

A. Appendix

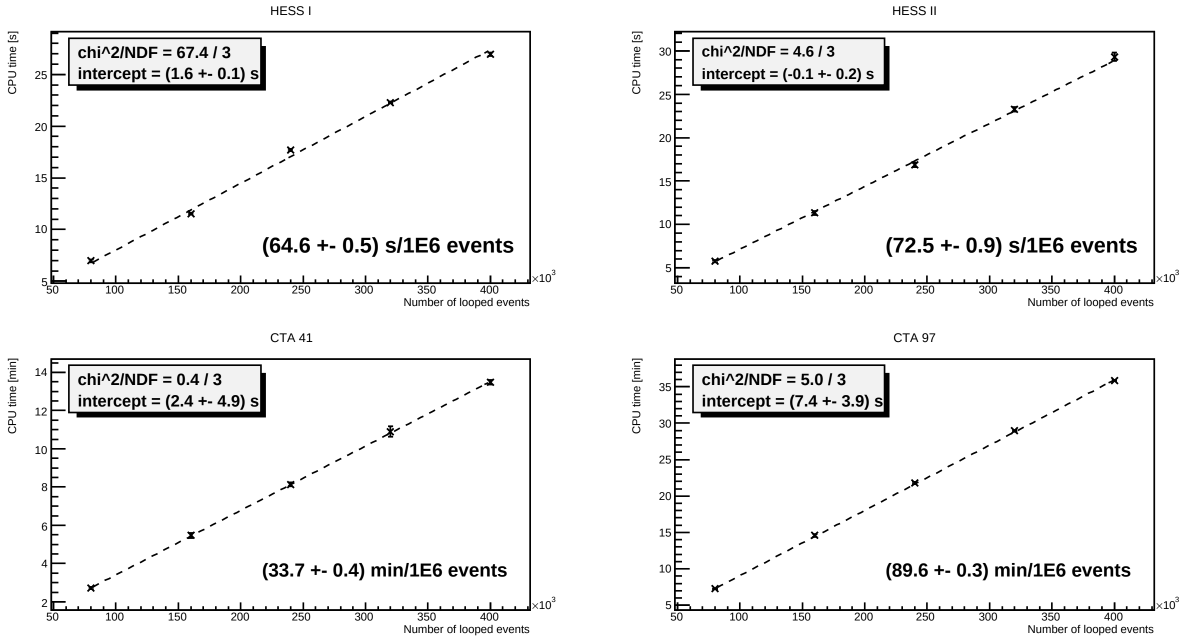

Fig. 8.— For one type of telescope array the CPU time to get the McTrueShower energy of one million events scales linearly with the number of looped events. The scale changed from seconds to minutes from the above to the below plot.