OpenDAFF

A free, open-source software package for directional audio data

Frank Wefers

Institut für Technische Akustik (ITA), RWTH Aachen,
E-Mail: fwe@akustik.rwth-aachen.de

Abstract

The *Open Directional Audio File Format* (*OpenDAFF*) is a new free and open-source software package for directional audio data, like directivities of sound sources (e.g. loudspeakers, musical instruments) and sound receivers (e.g. microphones, head-related impulse responses/transfer functions, HRIRs/HRTFs). OpenDAFF enables the exchange, representation and efficient of such data in form of single DAFF files (*.DAFF).

Under the term '*software package*' OpenDAFF unites the DAFF file format, a C++ software library for reading such directional audio files, as well as Matlab scripts for their creation and reading. The package is complemented by helpful tools, e.g. a graphical viewer for DAFF directional audio data, (called *DAFFViewer*) and a command-line utility (called *DAFFTool*).

The objective of the OpenDAFF project is to create a free and open-source format, that aims to ease and facilitate the interchange of directional audio data. OpenDAFF has been developed at the Institute of Technical Acoustics (ITA) at RWTH University, Germany in the context of real-time auralization. Consequently, are a minimal computational overhead as well as small memory footprints major design principles.

In this paper we introduce the OpenDAFF software package, describe its fundamental ideas and design principles and show up its functions.

Introduction

By the term '*directional audio data*' we refer to acoustic measures which depend on direction. Various examples for such data can be found in acoustics: The directivities of sound receivers (microphones, artificial heads, etc.) and sound emitters (loudspeakers, musical instruments, etc.), but also surface parameters like diffusion coefficients. Mathematically such measures can be expressed as functions defined on the surface of the sphere. All of the applications mentioned here are fundamental input parameters for acoustic computer simulations and auralization.

Usually the data itself is spatially sampled using a *sphere grid* and values are only known at discrete directions or points on the sphere surface. Each grid point is associated with a *record*, whose domain depends on the application and content. Two main groups of domains can be determined: a) signal-/system-related measures, like impulse responses (time-domain) and transfer functions (frequency-domain) and b) energy-related measures, defined for certain frequency bands, e.g. a magnitude spectrum in third-octave resolution covering the audible range.

In practice, sharing such directional data sets among colleagues or even across institute/company borders is usually hindered by a number of problems: Thinking of measurements, it is common to store them into seperate files - one for each direction. This usually results in a vast number of files which is inconvenient to handle. Individual files are usually tagged with their direction, by adding the angles into the filenames (e.g. HRIR_V123_H-32.wav). Other users might not be able to interpret these directions correctly: It might be unclear which spherical coordinate system was used and which handedness for rotations/angles was considered. Furthermore filenames cannot specify how the objection orientations relate to the spherical coordinates. This shall be illustrated this by two examples: Very often equiangular sphere samplings in lattitudes and longitudes (as found on the globe) are used. Many researchers perform HRIR/HRTF measurements with the vertical axis of the head matching the pole axis, meaning that the poles are above/below the head. Others align the pole axis along the ear axis. For speaker measurements it became popular to align the frontal axis of the speaker in direction of the poles.

All of these issues make the exchange of directional audio data complicated. Moreover they are a source of errors, which are hard to detect.

Motivation

We introduce OpenDAFF in order to avoid these problems by defining a common *'language'* for directional audio data. Thereby we want to ease the exchange of directional audio data between people and institutions. In our mind this can only be achieved by defining a universal semantic.

The objectives of the OpenDAFF project are:

- Design of a common interchange format for directional audio data
- Unambiguous, universal semantics
- Optimized for speed (real-time applicability)
- Multi-platform availability (Windows, Unix, MacOS, etc.)
- Free and open-source software (LGPL license)

The OpenDAFF software package specifies the DAFF file format and includes readers/writers and tools for this format. Currently the languages C++ and Matlab are supported.

Technology

The fundamental idea is to store the whole directional information, so basically all records among a spherical grid, within a single DAFF file. In order to meet a broad range of applications, OpenDAFF supports several types of content:

1. Impulse responses (IR)
2. Discrete Fourier spectra (transfer functions)
3. Magnitude/phase spectra (discrete frequencies)

Types 1-2 relate to signals/systems and are defined over an *equally-spaced* support (time or frequency). Category 3 refers to energy-related measures, which are defined for certain frequency bands, over *discrete* and not necessarily equally-spaced support. Aside this, data can be defined for an arbitrary number of individual *channels*. Thinking of an HRTF, there are to channels for each record - one for the left and one right ear.

The record data for a channel is an array of a fixed number of elements with a defined *datatype*. Impulse responses consists of real-valued elements, whereas DFT spectra have complex-valued coefficients. Each element can be quantized using 16-/24-bit signed integers or 32-bit floating points. For category 3 we restricted the quantization to 32-bit floating points only, because these data is usually very small (less than 1MB) which does not justify the effort of quantization.

The file format is designed to be very fast but also space-efficient. It is a container format, that arranges data within blocks (e.g. header, metadata, record data). A single file is loaded with very few I/O calls and the data alignment (memory boundary) is already cared for within the files. Concering impulse response content, we introduce a new technique we call *zero compression* and only store those filter coefficients with are effective, meaning they are unequal zero. Using this method we can reduce the file size significantly. A third-octave band magnitude directivity sampled in 5°x5° resolution takes only 350 kB. A HRIR measured with 1024 taps at 44.1 kHz and sampled with the same resolution results in a DAFF file of 4.8 MB.

Currently OpenDAFF supports only equiangular sphere samplings in lattitudes and longitudes. This is reasoned by its real-time background and the fact that this grid can be nearest-neighbour searched in constant time $O(1)$. For arbitary sphere grids the search is more time-comsuming. Support for these grids is planned for the future.

Common semantics

In order to solve the ambiguities with object orientations vs. data points, OpenDAFF introduces two separated views which are shown in figure 1: *Data spherical coordinates* (*DSC*) and *object spherical coordinates* (*OSC*). DSC are the reference frame for directions of records and have no relation to the object itself and we name the angles α *(alpha, A)* and β *(beta, B)*. A0° is the prime merdian (red) and B0° marks the south pole (green). OSC directions refer to the *object*: They are named *azimuth* (phi, φ, P) and *elevation* (theta, ϑ, T). We define that (P0°,T0°) points to the front, (P90°,T0°) left and T90° points up. Both views are mapped to each other by 3-D rotations (yaw-pitch-roll convention) within spherical coordinates. This *orientation* has to be determined just once and it is written into the DAFF file when it is created. The angle transformation is lightweight and fast. Users will prefer the OSC, because it unambiguously defines directions with respect to the object. Fallback on the DSC is only necessary for interpolation.

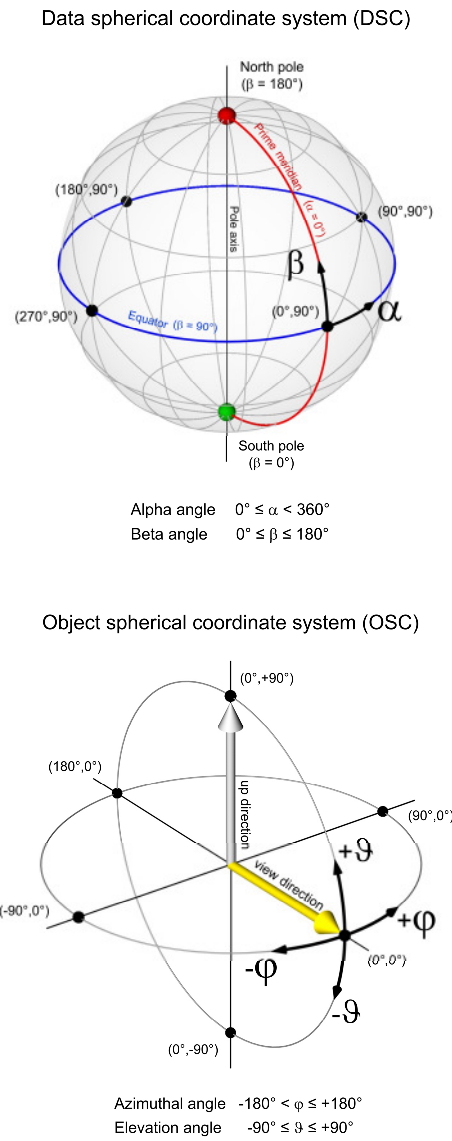

Figure 1: Separated data and object spherical coordinate systems used in OpenDAFF

Status and future

OpenDAFF is an open project. It is free software, published under the terms of the GNU Lesser General Public License (LGPL). Currently (March 2010) the project is in an early state. Most features already work. However we entitle the current state as a *draft*. Starting from here, the idea is to include expert commitees and to refine the current work. The outcome of this work shall be an official OpenDAFF specification.

For more information we refer to the project website: **http://www.opendaff.org**