ORCHIDS : Abstract and temporal orchestration software

End user documentation

IRCAM
1, place Igor Stravinsky
F-75004 Paris

ircam
Centre
Pompidou

This report was prepared by
Philippe Esling
Antoine Bouchereau

Supervisors
IRCAM

Philippe ESLING
Equipe Representations Musicales
Institut de Recherche et Coordination Acoustique / Musique (IRCAM)
1, Place Igor Stravinsky
F-75004 Paris
France

www.ircam.fr
Tel: (+33) 6 32 58 91 08
E-mail: esling@ircam.fr

Release date:	11 / 2014
Category:	1 (public)
Edition:	First
Comments:	This report is an internal technical report for IRCAM.

Contents

Chapter 1

Installation

1.1 Installation procedure

1. Open the disk image Orchids.dmg
2. Drag & drop the Orchids.app in the Applications folder
3. Download the required fonts for the Bach interface
4. Unzipped the sound database distribution (*SOL_0.9*)
5. Launch the Orchids.app in your Applications folder
6. Click on the *Database* tab on the top toolbar
7. Drag & drop your *SOL_0.9* folder in the explorer **on the pre-existing *SOL_0.9* icon**.

1.2 Contents of the package

The installation package is split to provide the following :

1. Orchids application
2. Orchids sound database

The sound database is a separate distribution which entails *Studio On Line* (SOL) and IRCAM's proprietary sound databases.

1.3 Troubleshooting

If you encounter any problem or the application doesn't seem to work, please report your problems in the user group devoted to this software:
http://forumnet.ircam.fr/user-groups/orchids/
You can also use this group to share your ideas with other users, and we encourage you to send us your good results to be featured in the front page of the webpage.
For other questions, you can send a mail directly to <*esling@ircam.fr*>, <*antoine.bouchereau@ircam.fr*> and <*danieleghisi@gmail.com*>.

Chapter 2

Main modules

2.1 General view

When launching Orchids, the main panel (cf. Figure 2.1) appears to the screen which contains

- The ***toolbar*** (cf. Section 2.2) allows to quickly access different features of the application.
- The ***activity monitor*** which displays information about currently running processes.
- The ***target information*** (cf. Section 2.3) which allows interaction with the current target.

After launching the *Orchids* application, it will automatically launch the database handler and start establishing the communication protocol with it. The whole software being constructed around a modular architecture, when the communication have been established, the activity monitor display "Ready" in order to indicate that the application is running correctly.

2.2 Toolbar

The toolbar can be found on top of the main panel. It is designed to allow quick and easy access to different parts of the software. Figure 2.2 shows a snapshot of this toolbar in the application. We quickly explain the functionalities related to each button of the toolbar in this section and provide pointers to more detailled explanations inside this document.

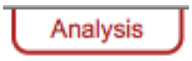

Opens the analysis panel which gives access to target features inspection and also criterias for the orchestration (cf. Chapter 2.3).

The explorer panel allows to browse the sound database and to perform queries on its content (cf. Section 3).

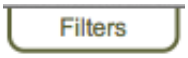

The filters allow to reduce the search space by filtering sounds to be used in orchestration based on their symbolic or spectral attributes (cf. Section 2.5).

The orchestra panel allows to define the instrumental staff and spatial positions that will be used for the orchestration search procedure (cf. Section 2.6).

The solutions panel provide controls over the orchestration (launching the search) and displays the orchestral solutions (cf. Section 2.8).

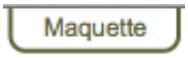

The maquette object allows to lay macro-temporal successions of orchestral solutions and provides a score editor (cf. Section 2.9).

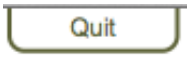

Close the orchestration software properly.

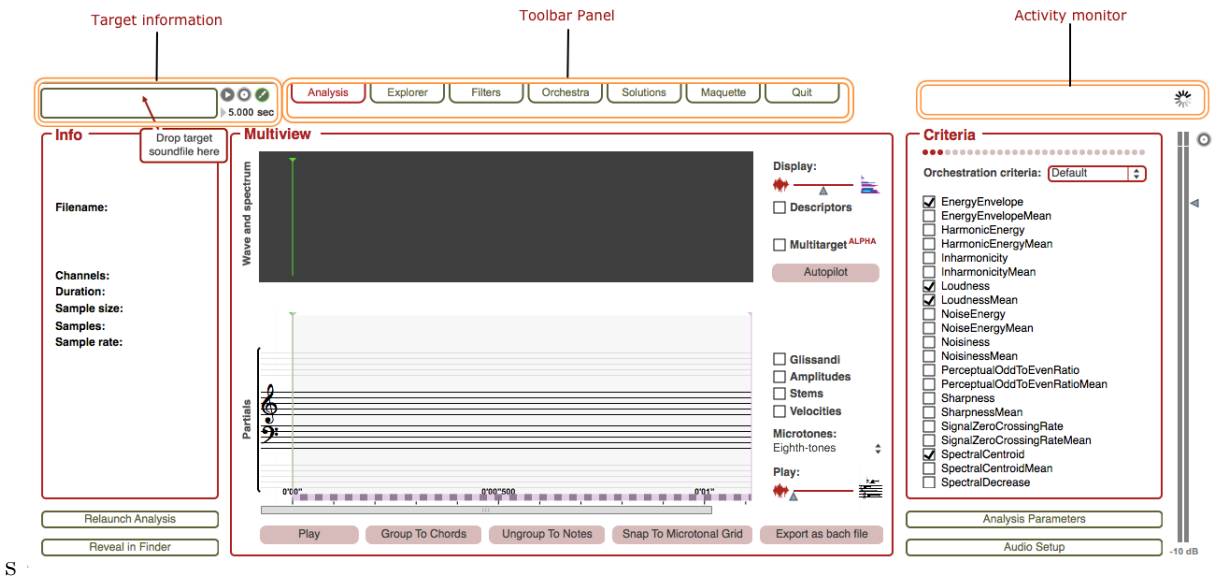

Figure 2.1: First minimal configuration of the application upon launching. It contains the *toolbar* panel (cf. Section 2.2), the *activity monitor* and the *target information* panel (cf. Section 2.3) which gives current target informations.

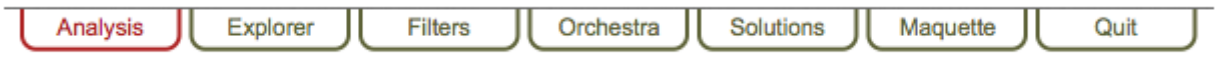

Figure 2.2: Toolbar summary for quick access to different functionalities of the software.

2.3 Target

For defining an orchestration problem, the target is the goal to be approximated as closely as possible by the algorithm. This search procedure is based on approximating several timbre properties jointly. The advantages of the Orchids system is that this approximation can be made separately over temporal shapes, mean values or multivariate distributions (Partials and Mel bands). Furthermore, for the definition of the target (and therefore the features themselves) several possibilities are given to the user.

- Using a single target sound file and selecting its properties
- Using some features from a first target sound file and then, thanks to the ***freezing*** mode, use other features from another sound file.
- Modifying any features shapes and values from a target and possibly mix it with original features.
- Starting from a void set (***abstract*** target mode) and directly drawing shapes to optimize.
- Finally the **multi-target** mode, which will perform several search sequentially but which is still in a very *alpha* primitive mode.

Of course all these modes can be mixed together depending on the specific needs of a user. The target panel is therefore divided into four different parts. (cf. Figure 2.3)

Info Informations about the target sound file.

Wave/Spectrum Display of the sound waveform and spectrum. The slider on the right allows to change the opacity of curves. When the Descriptors checkbox is ON, the display changes to a descriptor view that contains the temporal shape, the mean and standard deviation of each descriptor. A drop down menu allows to choose which descriptor to show and its temporal shape that is editable.

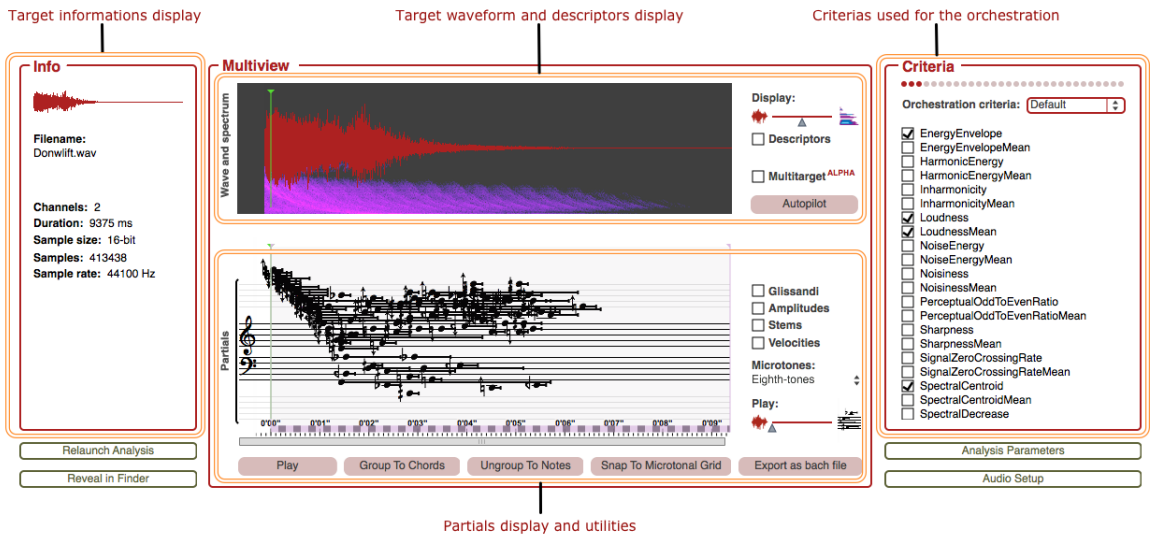

Figure 2.3: The target panel allows to set orchestration objectives either through multiple sound files or by directly drawing desired temporal shapes (*abstract* mode).

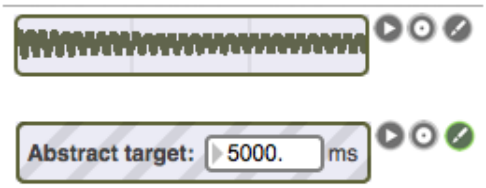

Figure 2.4: This target panel allows to set the taget to abstract mode or to record a new target.

Partials Display of the partial tracking of the target. Several display options are available on the right. The Play slider affects the listening of the target (click on Play) by creating a balance between the waveform and its partial tracking synthesis.

Criteria This is the list of all different criterias that can be used for the orchestration. Some presets are already available out-of-the-box, but you can save and load your own choices at any time.

Besides this main target panel, a small target controller (on the top left corner) can be accessed from every panel. It displays the waveform, allows to listen to the target, record a new one or to switch to the abtract target mode. (cf. Figure 2.4)

Abstract The abstract mode allows to define a target without the need to provide a sound file. By clicking this button, the system will provide all spectral descriptors corresponding to a second of silence. Starting from there, the user can start drawing new shapes to orchestrate a purely abstract target.

Duration Define the target length (in milliseconds) for optimization. This allows another level of flexibility where the original spectral features can be spread (or shrunk) over the timeline. This is especially useful for abstract targets where there is a need to impose a sense of global temporality.

2.3.1 Selecting a part of the target

Sometimes, it might be desirable to restrain the search to only a part of your current target (only a segment or one analysis window, especially in the case of static orchestrations). In that case, you can define the start and end points that you wish to orchestrate inside a target of any length or complexity.

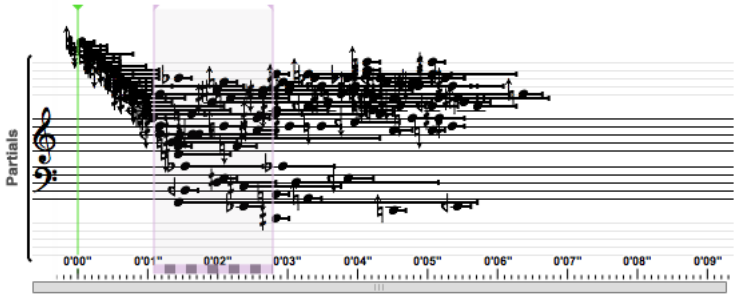

Figure 2.5: The selection window can be placed over the partial tracking in order to restrict the orchestration search only to a smaller portion of the target.

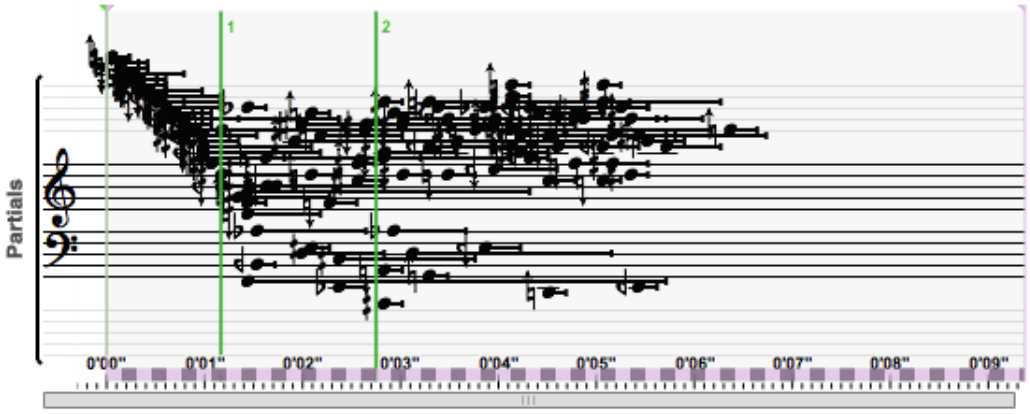

Figure 2.6: When using the multi-target mode, you can lay a set of markers. Each portion of the target will correspond to a separate search problem.

The search procedure will then only account for this part of the target. This can be done by simple dragging the purple selection frame inside the partial tracking window, as displayed in Figure 2.5.

2.3.2 Multi-target mode

Warning: This mode is in a very alpha state and can take a long time to compute.
By setting the search to the multi-target mode, it is possible to lay a set of markers (as displayed in Figure 2.6) and, hence, segment the target into multiple search procedure. Each portion of the target divided by the green markers placed after ticking the *multi-target* checkbox will be considered by the system as an entirely different search problem. The results of every search problems will be displayed on a common score in the Solutions panel.

2.3.3 Auto-pilot mode

If you just started using the orchestration and feel kind of lost inside the plethora of parameters, we devised the "Auto-pilot" mode specifically to help you start finding interesting orchestrations right away. This mode helps you to find the optimal set of criterias and algorithm parameters for a particular target. After putting a target inside the system, just click the "Auto-pilot" button and you will see that the set of criteria changed, as well as the algorithm selected. The auto-pilot might also detect that you are trying to input a target containing multiple events. It will then propose you a segmentation through a set of markers, and automatically set the search mode to "multi-target".

2.4 Features

The features panels allow you to select, modify or input entirely novel spectral feature objectives. The corresponding choice can either be to optimize the temporal shape (normalized in the range [-1,

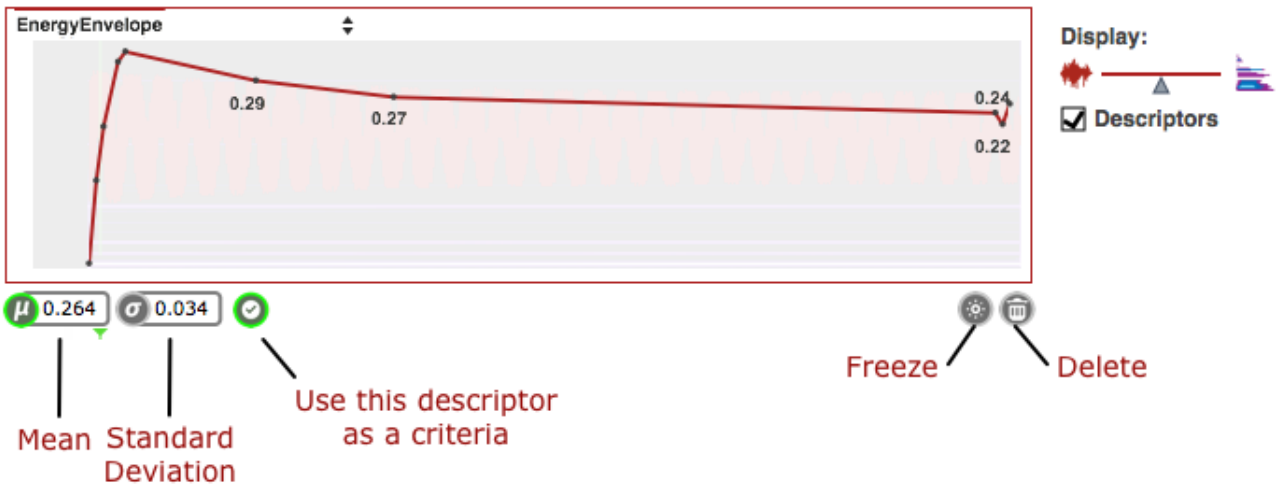

Figure 2.7: Features panels allow to select, modify or input objectives. The corresponding choice can either be to optimize the temporal shape (normalized in the range [-1, 1]), optimize the mean value or optimize the standard deviation. It is also possible to optimize any combinations of these three separately.

1]), optimize the mean value or optimize the standard deviation. It is also possible to optimize any combinations of these three separately. Figure 2.7 summarizes the action available from the features panel.
Selecting the checkbox on the right indicates that the display switch to the Feature mode. For each shape, the curve is editable and the mean and standard deviation are automatically updated.

Delete Erase all time points and empty the current temporal shape. This can allow you to direclty draw your desired curve of descriptor to optimize.

Freeze Keep this descriptor shape, even if a new target is input to the system. This will allow you to mix the features of different sound targets.

2.5 Filters

Filters are a complete class of constraints which can be imposed on the search procedure. The idea is to reduce the search space to use only certain types of sounds by selecting only the values that should be used throughout the search. The filters panel (cf. Figure 2.8) is divided into the symbolic filters (left) and the spectral filters (right). This can allow you for instance to try to orchestrate only with *pizzicati* sounds or sounds that have their *mean EnergyEnvelope* between a desired range.
Symbolic filters are used to constrain the sounds used based on their symbolic descriptions. (eg. *playingMode*, *note*, *dynamics*). For example, the orchestration solutions can be constrained to use only *mezzoforte* sounds to approximate the target. Each checkbox can be changed independently. Once the modifications are made, the filter must be applied to the server
Spectral filters are used to constrain the sounds used based on the values of their spectral features. For example, the orchestration solutions can be constrained to use only sounds which have a very high deviation in energy (eg. *cresendo*, *decrescendo*, *tremolo* and so forth) to approximate the target. The slider (cf. Figure 2.10) allows to directly impose a range of allowed values. *Once the modifications are made, the filter must be applied to the server.*

2.6 Orchestra

The orchestra panel allows to define which instruments should be used in the orchestration search and the numbers of each instrument in the orchestra. The top panel (cf. Figure 2.11) provides a preset

Figure 2.8: Filters allows to constrain the search space by selecting only the values that should be used in the search. The filters panel is divided into the symbolic filters (left) and the spectral filters (right)

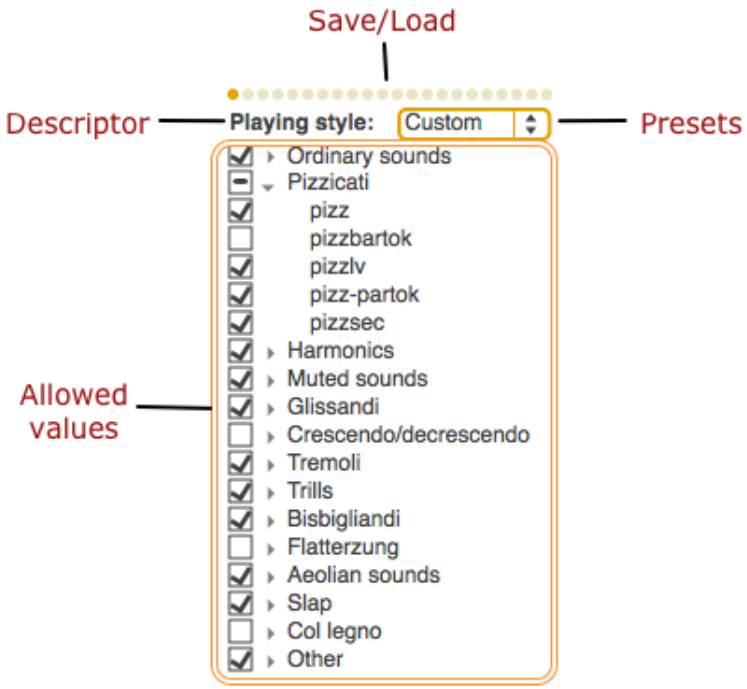

Figure 2.9: Symbolic filters are used to constrain the sounds used based on their symbolic descriptions. The top buttons allow to directly select or deselect all values. Each checkbox can be changed independently. *Once the modifications are made, the filter must be applied to the server*

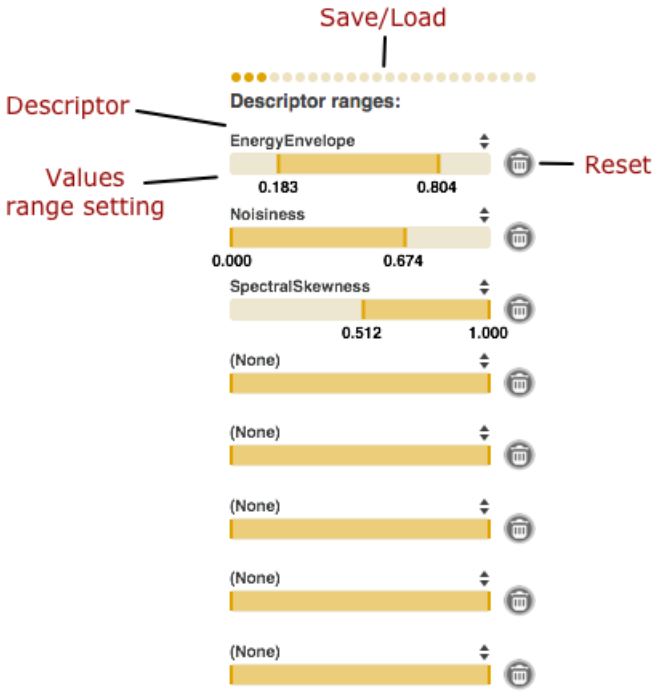

Figure 2.10: Spectral filters are used to constrain the sounds used based on the values of their spectral features. The slider allows to directly impose a range of allowed values. Once the modifications are made, the filter must be applied to the server.

system which allows to save and load orchestra configurations for faster handling. The interesting part to understand is that the orchestra is defined through its players mainly. Therefore, one player can be assigned to multiple instruments (that will therefore not play at the same time). The user should define an instrumental configuration and how many players follow this configuration and then add them. After adding players, independent remove buttons are available next to each player.

The stage mixer allows to change the spatial position of each player (by moving it in the stage) and to control the gain of each instrument (by pressing *Ctrl+click & drag* on the instrument).

The settings panel allows to change the *microtonal resolution* for the orchestration (*we usually recommend to use quartertones resolution for a better solution approximation*) and the spacing reverb of the stage. *Once the modifications are made on the interface, the changes must be applied to the server (bottom button)*

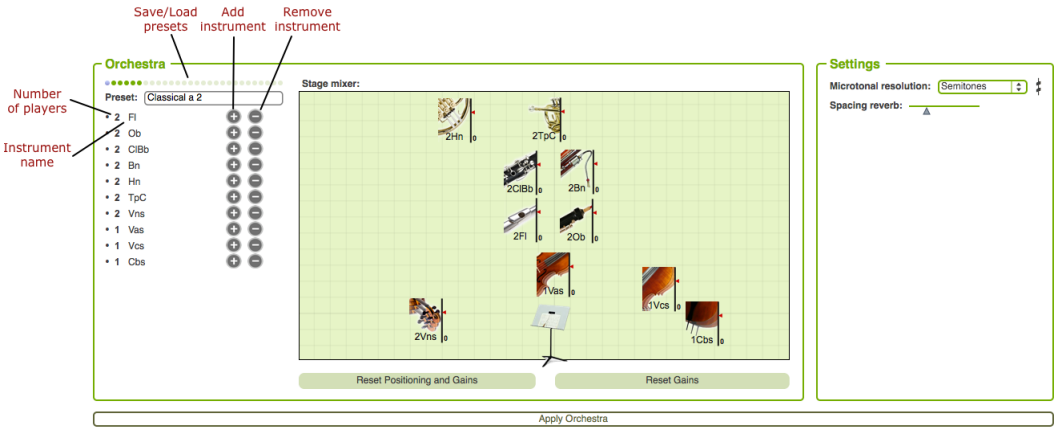

Figure 2.11: The orchestra panel allows to define the players and their instruments. The top part also provides a preset system to load and save orchestra configurations.

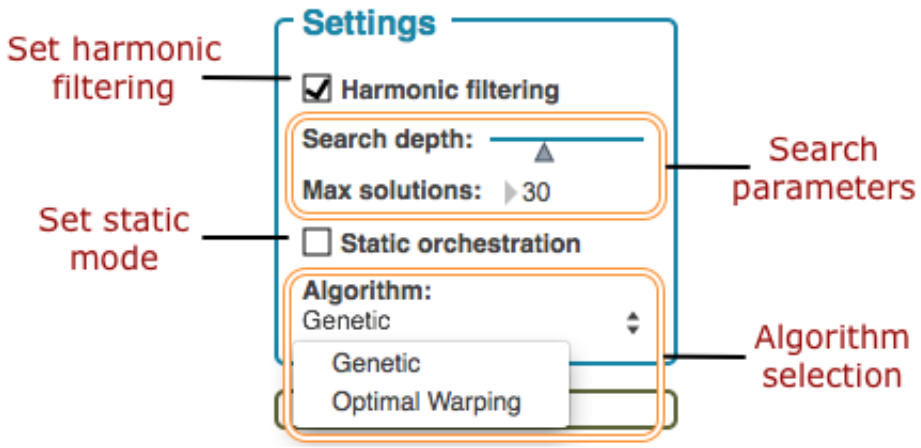

Figure 2.12: The algorithm panel allows to select which search algorithm should be used as well as its internal parameters. Currently two algorithms are available, but more may be coming thanks to the modular architecture of this software.

2.7 Algorithms & Settings

The modular structure of the *Orchids* system (cf. Figure 2.12) allows for several completely different search algorithms and parameters. That way, the user can select its preference and best approach for each particular target. The algorithm panel in the Solutions panel allows to select which search algorithm should be used as well as its internal parameters. Currently two algorithms are available, but more may be coming. The *genetic* algorithm is based on hypervolume-based genetic evolution and the *Optimal Warping* algorithm is based on a combination of multi-objective and time series matching algorithms. The optimal warping search is slightly slower than the previous one but is a lot more fit to temporal targets.

Harmonic Harmonic filtering allows to restrain the search space by automatically deriving the note used from the partials analysis (it is recommended to keep this setting, but removing it can lead to more search space possibilities).

Static The static mode restrains the search by using only mean features criterias and onsets to zero (all instruments will start at the same time).

Depth The search depth corresponds to the number of iterations in the search process.

Algorithm Selects between *Genetic* (recommended for static targets) and *Optimal Warping* (recommended for temporal targets).

2.7.1 Advanced search parameters

All search parameters are now available by clicking on the "Advanced search parameters" buttons in the "Solutions" panel. You can select various operators and change the size of populations in order to find a good tradeoff between speed of the search and quality of the solutions. Note that each operator enhance the quality but some might take a lot of time (especially the *local optimization*). Furthermore, wider population sizes will increase the quality of solutions but greatly reduce the speed.

Selective-Mutation Perform a loudness-sensitive mutation of individuals

Harmonic-Fix Perform a solution-specific fixing of the harmonic distribution. Its occurence is controlled by the frequency setting *(the lower the number, the more this operator will be applied).* This operator is only effective with PartialsMeanAmplitude.

Local-Optimization Hill-climbing operator that works by locally searching around Pareto-optimal solution (warning, this is a memory-intensive operator)

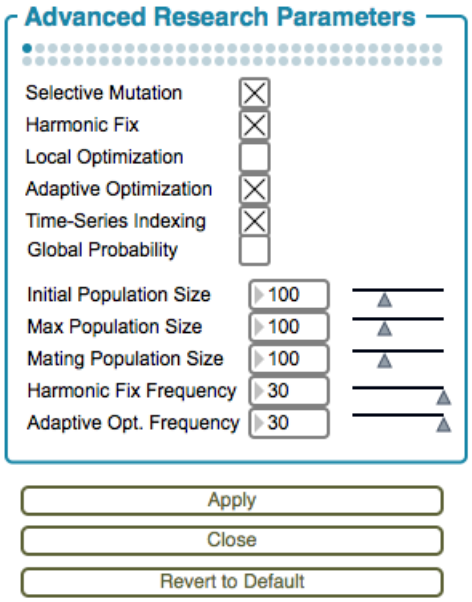

Figure 2.13: The advanced search parameters panel.

Adaptive-Optimization Perform an adaptive optimization of solutions (the lower the number, the more this operator will be applied). This operator is only effective with PartialsAmplitude or a temporal criteria and the *OptimalWarping* algorithm.

Time-Series-Indexing Rely on time series indexing to enhance the current population. This operator is only effective with Adaptive-Optimization, OptimalWarping and temporal criteria.

Global-Probability Probabilities of crossover and mutation are globally adapted along the search to match the current state of the system.

2.8 Solutions

Once the search procedure has been launched by the user, *Orchids* will start computation. Meanwhile, the activity monitor displays what operations are being performed, that signify that the search is undergoing. This panel allows to parse through the *current generation* of solutions. As the algorithms are based on evolutionary procedures, several generation will be created and the client automatically updates the solutions based on the advances of the server. The results panel (cf. Figure 2.14) is divided in four distinct parts.

The left panel (cf. Figure 2.15) allows to see the position of the solutions on the optimization map (and therefore how solutions approximate different objectives more or less). Around this map, the user can select two of the spectral descriptors. Finally, the dots on the map are clickable and will automatically display the corresponding solution in the score part.

The middle part of this panel is a symbolic score representation based on the *Bach* software. It allows to directly see the score corresponding to a solution and (if the sound database is installed on the hard drive) to play the sound mixture straightforwardly. One of the main advantages of this representation is that it can be used in conjunction with the *Maquette* object (cf. Section 2.9) that allows to place several solutions on a longer time scale and therefore experiment with orchestral macro-articulations. It is therefore possible to copy, paste, add or modify any part of this representation. Furthermore this representation includes an automatic symbolic search system. This allows for example to change the pitch, duration or dynamics of any note and the system will find the corresponding sound file in the database. This can allow full transposition or rescaling of the solutions. The right part include interaction buttons to modify the representation, play the solution, save and load solutions, copy the current solution to the *Maquette* panel and export solution as audio, bach, MIDI, OpenMusic or PWGL file.

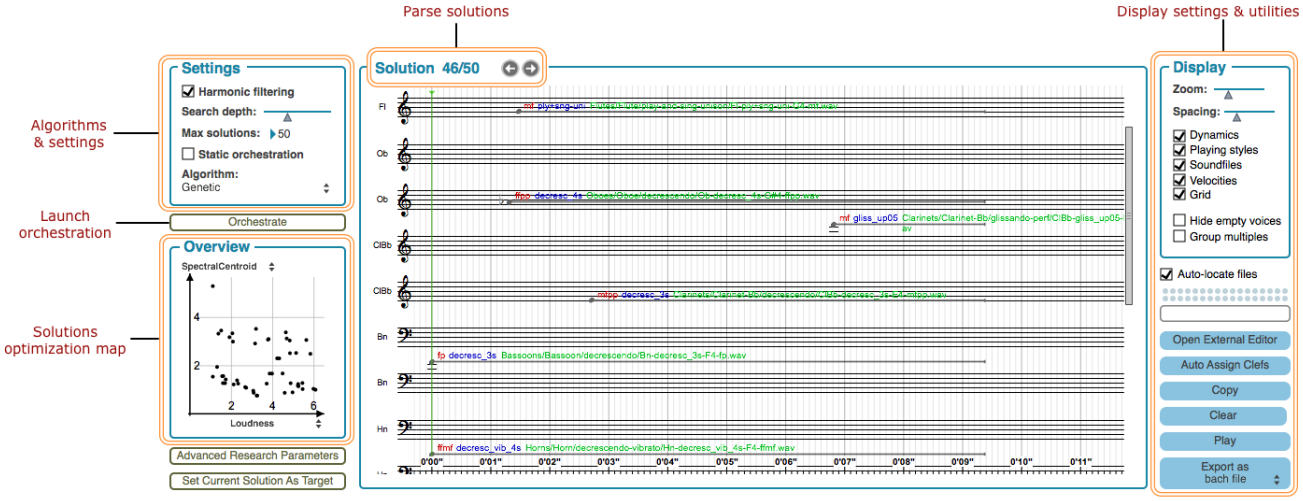

Figure 2.14: The results panel allows to parse through the orchestration proposals in real time.

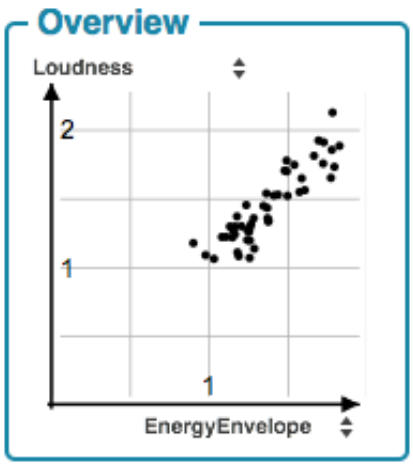

Figure 2.15: The left panel shows optimization information.

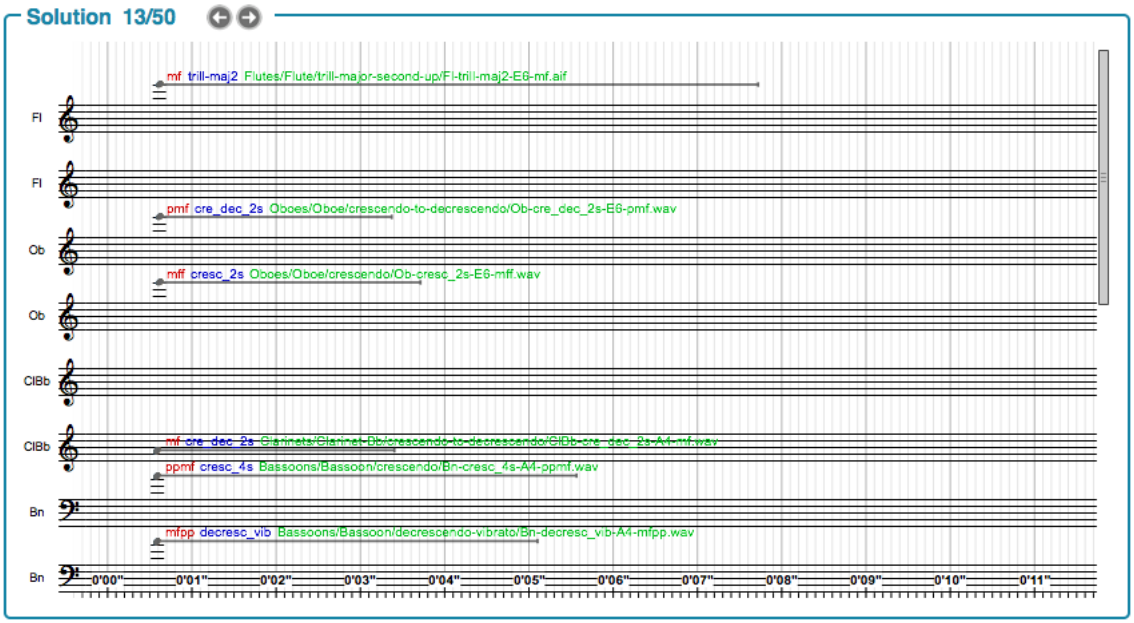

Figure 2.16: The score representation allows to directly see the symbolic representation of solutions. The right buttons can modify the aspect of this representation.

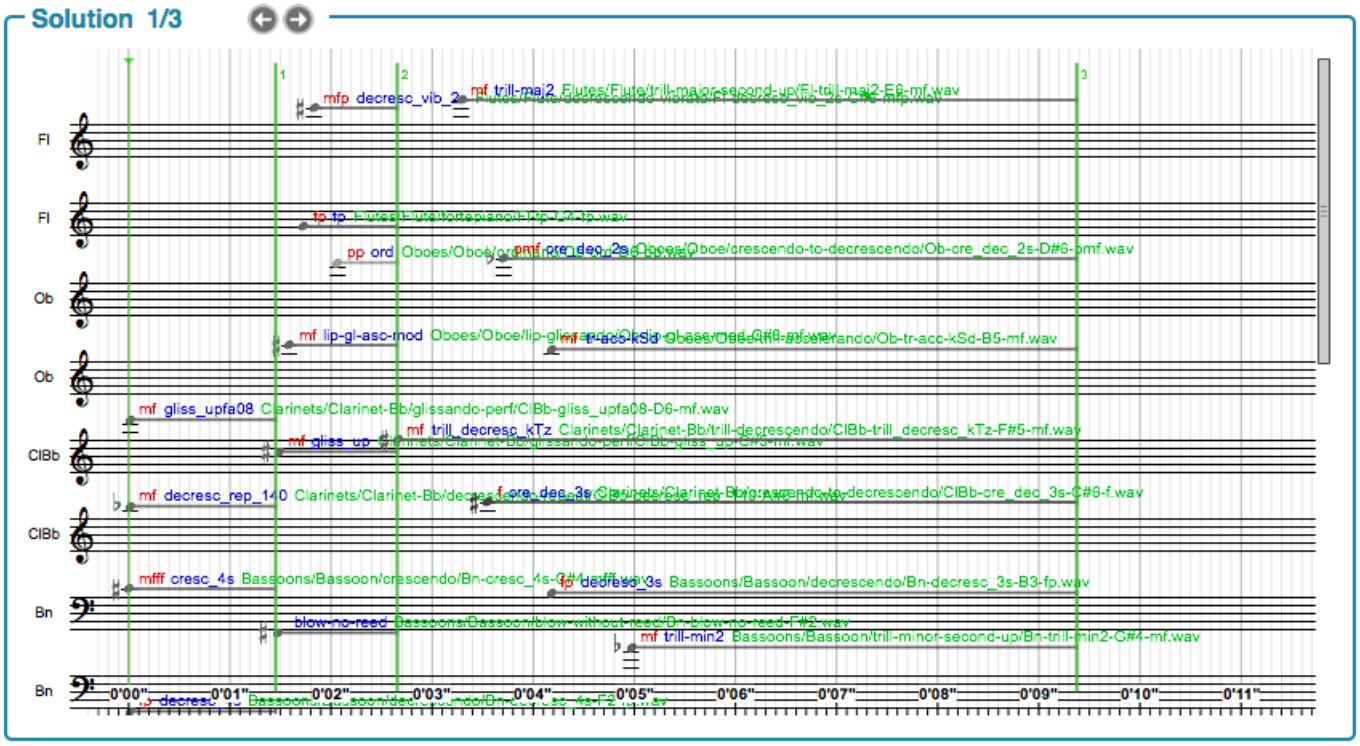

Figure 2.17: In multi-target mode, the solution panels will display the results of successive orchestration search along with the set of markers.

2.8.1 Multi-target solutions

Warning: This mode is in a very alpha state and can take a long time to compute.
When the multi-target mode is selected, and markers have been selected in the target panel, the solutions will be displayed as shown in Figure 2.17.

2.9 Maquette

When accessed through the toolbar, an extended score representation is opened. This timeline can be seen as a full sketch of potential orchestrations. As it is based on the *Bach* system, it embeds all the functionnalities offered by this system. The Maquette panel (cf. Figure 2.18) is divided in two different window.
The Quantize button creates a quantized version of the score than can be displayed by draging the bar from the bottom to the top of the window (“Drag me”) (cf. Figure 2.19).

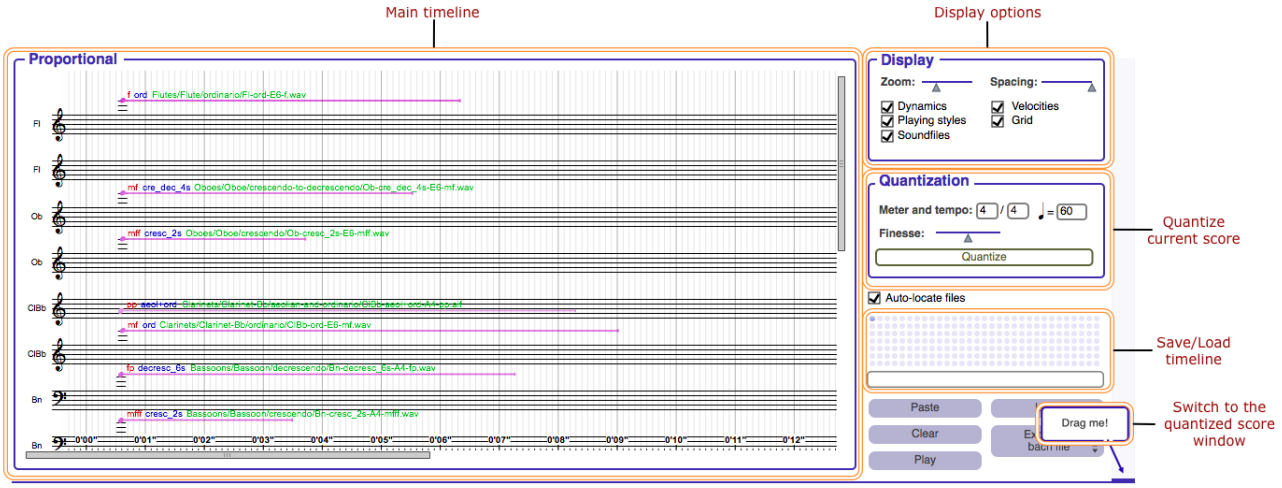

Figure 2.18: The maquette panel allows to construct a timeline of different orchestration solutions.

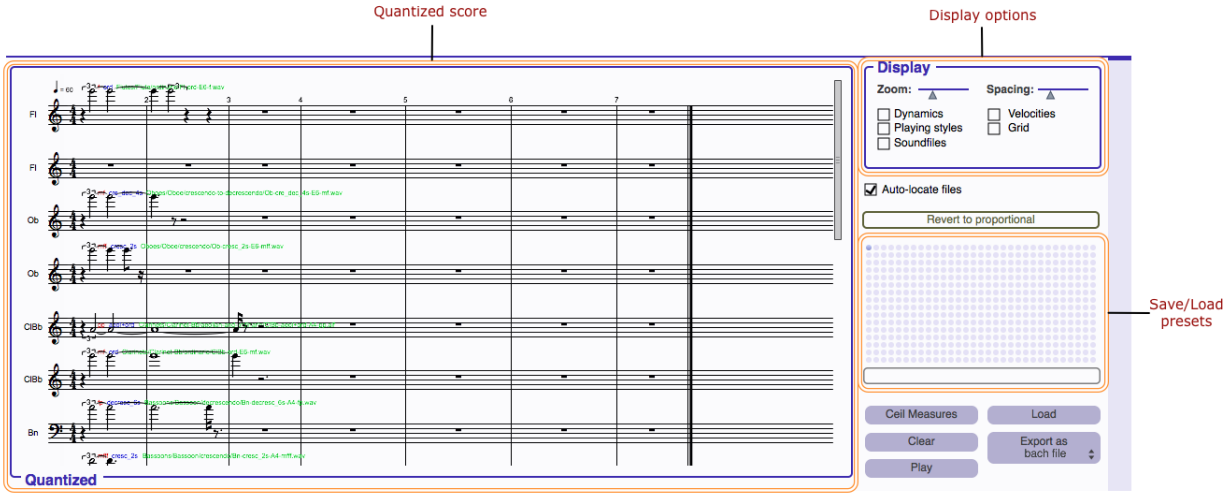

Figure 2.19: Quantized version of the timeline.

Chapter 3

Database querying

In addition to the orchestration search algorithms, the *Orchids* features an innovative "intelligent database" system. This system allows to perform queries for sound files based on the temporal evolution of their spectral descriptors. Therefore, when searching for specific properties of a sound, it is possible to directly draw the desired temporal shapes of a specific spectral descriptor and the system will retrieve the top matches from the database.

3.1 Introduction

From the toolbar, it is possible to access this system by clicking on the Explorer icon. This opens the first database panel (cf. Figure 3.1).

3.2 Explorer

The database explorer allows to go through the sound database by parsing the tree structure (cf. Figure 3.2).
By clicking on a sound sample, the Inspector panel displays symbolic informations, waveform and temporal descriptors of the sample that can also be listened (cf. Figure 3.3). The drop down menu allows to choose the descriptor to display.

3.3 Querying

We detail in this section each type of querying available

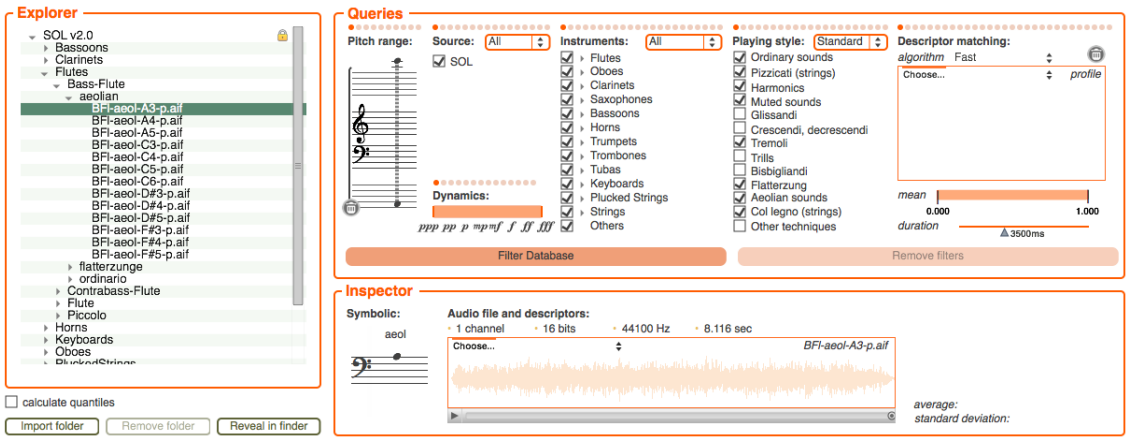

Figure 3.1: The database panel gives access to three different sub-systems for querying

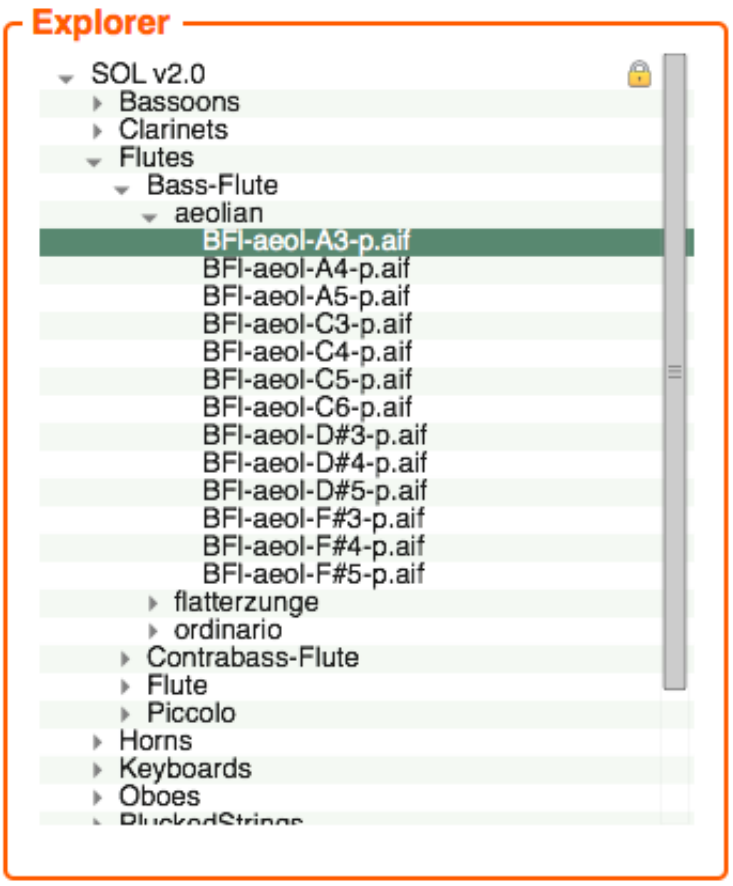

Figure 3.2: Exploring the sound database.

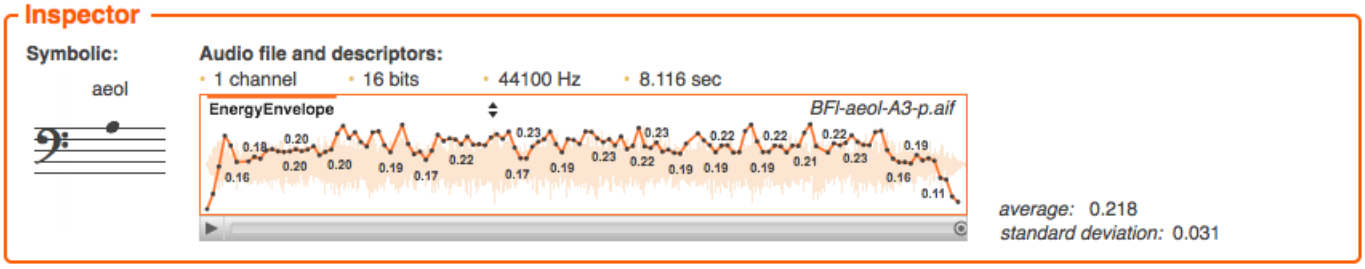

Figure 3.3: Exploring the sound database.

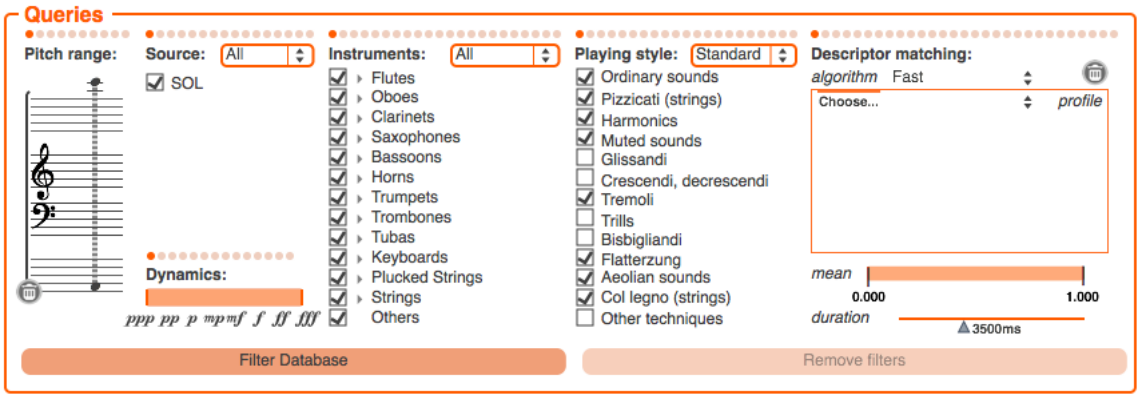

Figure 3.4: Values quering panel.

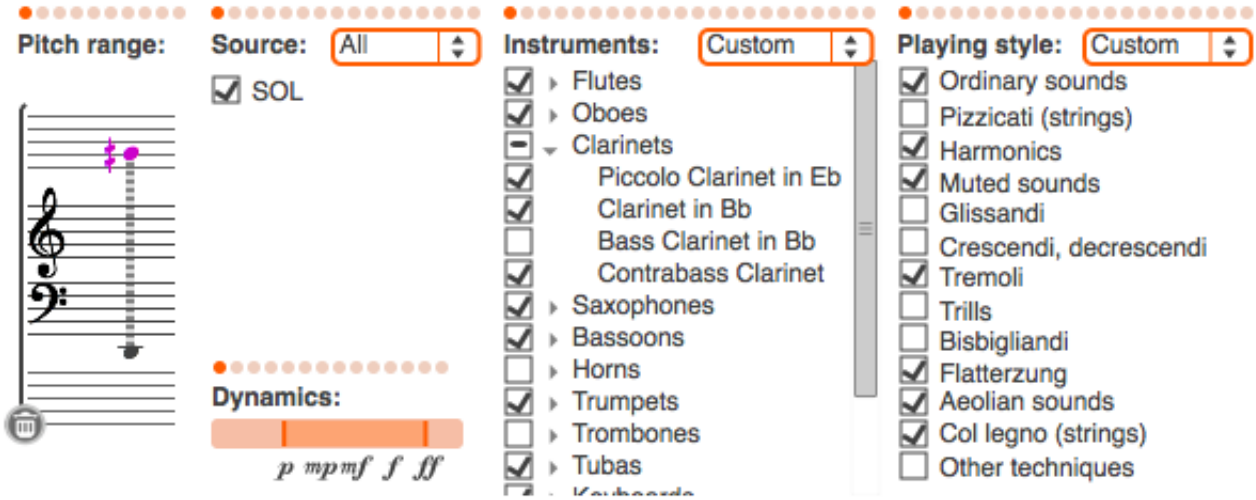

Figure 3.5: The samples querying panel allows to parse through all the samples in the sound database based on symbolic categories.

3.3.1 Values querying

The values querying panel is intended to have a quick lookout over every existing values in the database. The panel is divided in two distinct parts (cf. Figure 3.4).

3.3.2 Samples querying

The samples querying panel (cf. Figure 3.5) allows to find the samples based on their symbolic categories. Therefore, the first fields allow to select, the instrument, its note and playing style. Based on these informations, several samples are proposed in the corresponding field mostly based on different dynamics and variations of the sample. Finally, the left part of the main panel shows samples corresponding to the query.

3.3.3 Temporal querying

The temporal querying panel (cf. Figure 3.6) is based on an innovative "intelligent database" system. This system allows to perform queries for sound files based on the temporal evolution of their spectral descriptors. Therefore, when searching for specific properties of a sound, it is possible to directly draw the desired temporal shapes of a specific spectral descriptor and the system will retrieve the top matches from the database. The top field of this panel allows to select the desired spectral descriptor. The middle part is based on a breakpoint function and allows to draw any type of temporal evolution required. Finally the top right of this panel allows to erase the temporal shape or launch the query.

Complete Descriptor matching is computed by using a Dynamic Time Warping algorithm.

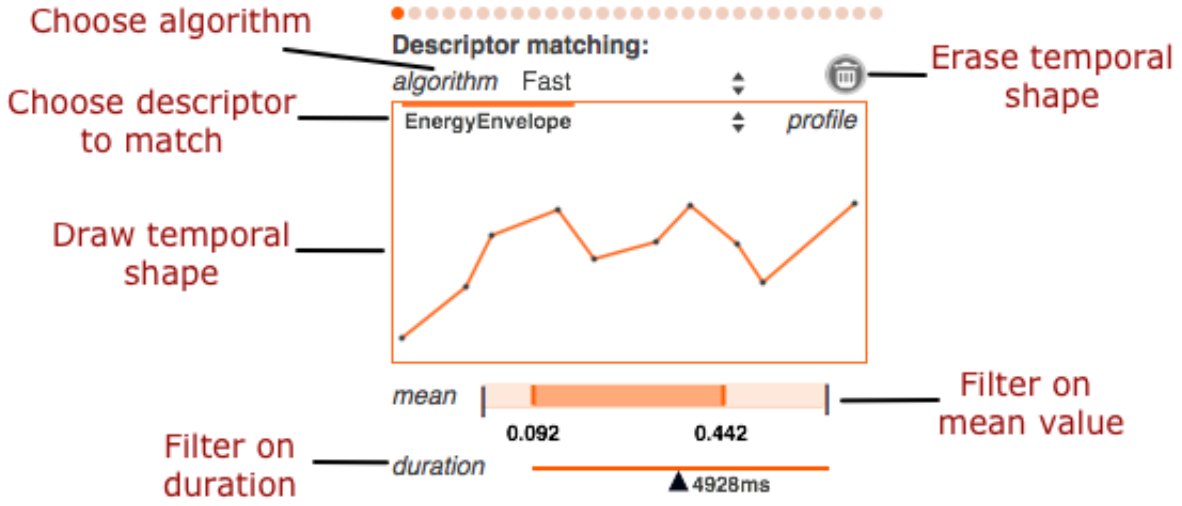

Figure 3.6: The temporal query panel allows to perform queries based on the temporal shapes of spectral descriptors.

Fast Descriptor matching is computed by a time-series indexing.

3.4 Results

When a temporal query is performed on the database, the Explorer panel (cf. Figure 3.2) automatically displays the matches of the query and their related spectral features are displayed by clicking on one sample in the Inspector panel (cf. Figure 3.3).
The results are the best in terms of time series distances but are shown in alphabetic order.

Chapter 4

Other features

4.1 Use Cases

This part is intended to explain step by step some new features of the *Orchids* system.

4.1.1 Add your own sounds

- Go to the Explorer panel
- Drag & drop your sound folder from the Finder to the database explorer (cf. Figure 3.2)
- A pop-up menu asks you to enter the sound folder instrument name (cf. Figure 4.1)
- The analysis of your sound folder begins, the activity monitor displays which sound file is being analyzed
- Your sound folder can be explored in the database explorer panel (cf. Figure 3.2)

4.1.2 Orchestrate with your own sounds

Warning on the naming of your sounds. If you want to ensure that your sounds will be used by the systems, for pitched sounds, you should first check if the sounds you added followed the required nomenclature. Indeed the system requires the sounds filename to be in the following form *<Instrument acronym>-<playing style>-<note>-<dynamics>.*
For instance, adding samples of a *Bassoon* (with the acronym *Bn*) playing in *ordinario* should all be in the same folder and named in the following way
Bn-ord-A4-mf.wav
Bn-ord-C#2-ff.wav
[...]

Figure 4.1: Enter the short instrument name.

For unpitched sounds, you can

- Go to the Orchestra panel
- Add a new instrument (cf. Figure 2.11)
- By clicking on the instrument name, your added instruments can be selected at the bottom of the drop down menu
- The user added sounds name may not have the same syntax than SOL so the symbolic notes may not be retrieved. It is recommanded to turn off the harmonic filtering (cf. Figure 2.12)

4.1.3 Removing sound folders

- Go to the Explorer panel
- Click on the folder you want to remove
- Click on the button "Remove Folder" at the bottom of the explorer

4.1.4 Add/Remove sounds from an existing folder

The new *Orchids* system allows you to directly modify your folders in the Finder, every modifications will be considered at the launch of the application.

- Go to your sound folder in the Finder
- Add or remove some sound files
- Launch Orchids
- The activity monitor displays which sounds are added or removed to the database

4.1.5 Freeze target features

- Go to the Analysis panel (cf. Figure 2.3)
- Click on the "Descriptors" checkbox, then the Descriptors panel displays (cf. Figure 2.7)
- Choose the descriptor you wants to freeze
- Click on the "Freeze" button
- Drag & drop a new target (or switch to the abstract target mode), the freezed descriptors are now available

4.1.6 Set current solution as Target

- Launch an orchestration
- Go to the Solutions Panel and choose one solution
- Click on the "Set Current Solution As Target" button
- The current solution is exported on wav file and automatically added and analyzed as a target

4.2 Bach commands

4.2.1 Basic commands: keyboard

- *Backspace*: Delete all selected items, except for selected note tails, for which the backspace key reverts their pitch to the original notehead pitch (no glissando).
- →: Shift selection onsets to the right; for selected tails: increase corresponding note duration
- ←: Shift selection onsets to the left; for selected tails: decrease corresponding note duration
- *Alt+*→: Increase duration for selected notes (possible breakpoint positions are rescaled). Add the Shift key to change duration more rapidly.
- *Alt+*←: Decrease duration for selected notes (possible breakpoint positions are rescaled). Add the Shift key to change duration more rapidly.
- *Alt+Ctrl+*→ (mac) or Alt+Shift+Ctrl+→ (win): Increase note duration, preserving the absolute position of each pitch breakpoint. Add the Shift key (mac) to change duration more rapidly.
- *Alt+Ctrl+*← (mac) or Alt+Shift+Ctrl+← (win): Decrease note duration, preserving the absolute position of each pitch breakpoint. Add the Shift key (mac) to change duration more rapidly.
- ↑: Raise pitch of selection by one step. Add the Shift key to raise the pitch by one octave. Add the Ctrl+Shift key to prevent notes from being assigned to different voices.
- ↓: Lower pitch of selection by one step. Add the Shift key to lower the pitch by one octave. Add the Ctrl+Shift key to prevent notes from being assigned to different voices.
- *Cmd+Z* (mac) or Ctrl+Z (win): Undo
- *Cmd+Shift+Z* (mac) or Ctrl+Y (win): Redo
- *Cmd+L* (mac) or *Ctrl+L* (win) when items are selected: Lock/unlock selected items
- *Cmd+G* (mac) or Ctrl+G (win): Group/ungroup selected items
- *Ctrl+A, Ctrl+B, Ctrl+C, Ctrl+D, Ctrl+E, Ctrl+F, Ctrl+G* (mac) or Shift+Ctrl+A, Shift+Ctrl+B, Shift+Ctrl+C, Shift+Ctrl+D, Shift+Ctrl+E, Shift+Ctrl+F, Shift+Ctrl+G (win): Transform pitches of selected notes into the corresponding pitch (A = La, B = Si, C = Do, and so on). The octave is chosen to be the one so that the snapped note is as near as possible to the previous note.

4.2.2 Basic commands: mouse

- *Cmd+Click* (mac) or Ctrl+click (win) on a notehead: Delete note
- *Cmd+Click* (mac) or Ctrl+click (win) elsewhere: Add a new chord having a note at the clicked position
- *Alt+click+drag* on a notehead: Copy the note into a new chord. Add the Shift key to add the new note to the existing chord
- *Click+drag* on a notehead or a duration line: Change the note pitch (dragging up/down) and the chord onset (dragging left/right). Add the Shift key to change either the note pitch or the chord onset. Add the Shift+Ctrl while dragging up/down to prevent notes from changing voices. Add the Shift+Cmd keys (mac) or the Shift+Ctrl keys (win) while dragging left/right to finely edit the chord onset.

- *Click+drag* on a note tail: Change the note duration (dragging left/right) and pitch of the note tail (dragging up/down provided that the allowglissandi attribute is set to 1). Possible pitch breakpoint positions are rescaled. Add the Shift key to change either the note duration or the note tail pitch. Add the Shift+Cmd keys (mac) or Shift+Ctrl keys (win) to fine edit the position. Add the Alt key to keep the existing absolute pitch breakpoint positions. This can also be combined with the previous Shift key addition
- *Doubleclick* on voice names, marker names or lyrics: Edit names directly
- *Doubleclick* on selected items: Selection off-line play (as the V key), but only if dblclicksendsvalues is set to 1
- *Shift+click+drag* on a clef (at the beginning of the staff): Move voice vertically (along with all the next ones), i.e.: modify vertical space before clicked voice
- *Shift+Alt+click+drag* on a clef (at the beginning of the staff): Move all following voices vertically (but not the clicked one), i.e.: modify vertical space after clicked voice
- *Right click* or two-fingers-tap: Open contextual popup menu. This can be opened by clicking on noteheads (note or selection popup menu), clefs (voice popup menu) or on the background (backgroun popup menu).

4.2.3 Others commands

To see all commands and shortcuts about bach.roll and bach.score, please refer to the Bach documentation.

Chapter 5

Troubleshooting

5.1 Common problems

Before going to a mind-blowing puzzle, always remember that the status console (top right corner) of the application can indicate the current status of the system. If this not helps, you can still check the message console by pressing *Cmd+M*.

5.1.1 The application seems to freeze at the beginning.

This can be normal depending on your OS version. The application needs to check the indexing tree for all your soundfiles which might take a few minutes. Give it a couple of time and the application should respond normally afterwards.

5.1.2 When I drop the SOL_0.9 folder, the application analyze it completely

This can happen if you did not drop the folder correctly. Please make sure that you drop the folder **on the pre-existing SOL_0.9 icon** (that should be crossed and red at the beginning, then becomes normal).

5.1.3 No solutions appear after I clicked 'orchestrate'

Apart from "real application crashes", several reasons can lead to an absence of orchestration.

1. Your search space is empty. This might come from trying to orchestrate for instance a low-pitched sound only with high-pitched instruments.

2. Filtering removed all sounds from the space. By providing "incoherent" configurations, you deprive the algorithm from any search possibility.

3. The search is not finished. Some configurations of criteria, orchestra and search parameters can take quite a long time depending on your computer. Make sure that the search is not still working by looking at the status panel and the message panel (*Cmd+M*). If you want to speed up the search, please read the next section.

5.1.4 I don't know how to use the parameters

If it is your first time using the software, you can try to check the Auto-pilot mode. If you want to see its effect, it will influence

- *Criteria.* The auto-pilot analyze which descriptors exhibit the *highest variability* and automatically selects these.

- *Target mode.* If your target shows no clear temporal evolution, the auto-pilot will select the *static* mode, if the target shows a clear temporal shape, it will select the *temporal* mode. If you tried to input a multi-events complex target, it will select the (**unstable**) *multi-target* mode.

- *Algorithm.* For static targets, the auto-pilot will select the *genetic* algorithm, and for temporal targets the *optimal warping* algorithm.

5.1.5 The system is very slow when orchestrating

The speed of the system is directly correlated to the complexity of the search problem you are optimizing. For instance *PartialsAmplitude* is the temporal version of *PartialsMeanAmplitude*, hence its computation will be as longer as the length of your target. If you want to speed up the search, several factors can be played with, but note that this might come with a net decrease in the quality of the solutions.

1. *Criteria.* The number of criteria influence the search performances, you can try with minimal sets of criteria (Partials, Loudness) to enhance the speed. Depending on the temporality of your target (static or temporal), you might prefer choosing static criteria (Partials*Mean*Amplitude, Mel*Mean*Amplitude and Loudness*Mean*), to improve the search speed.

2. *Filters.* Filtering your search space (reducing the number of possible playing styles, dynamics, etc...) will also increase the speed of the algorithm

3. *Orchestra.* Reducing the number of instruments, and allowing only semi-tones resolution will speed up the search

4. *Algorithms.* The *Genetic* algorithm is quite faster than the *OptimalWarping* algorithm.

5. *Advanced search parameters.*

   (a) All search operators (*selective mutation*, *harmonic fix*, *adaptive optimization*, *local optimization* and *TS index*) take their share of time but can be de-activated at will in the "*Advanced search parameters*" panel.

   (b) Lower population sizes and number of iterations will also speed up the search and can be changed in the "*Advanced search parameters*" panel.

5.1.6 The quality of my solutions are not good

Several points can lead to a variability in the quality of solutions.

1. *Your target.* Some targets can be impossible to reconstruct. One of the most common error is to try to orchestrate a "multi-event" target directly, with very different parts in the same sounds. You can try this with the *multi-target* mode, but you can also start by selecting a sub-section of your target. However, keep in mind that even some single events might be impossible to reconstruct.

2. *Search parameters.* By increasing the search parameters and giving more possibilities to the algorithm, you can increase the probability to find a closer solution. You can read the previous section and apply the opposite advices (reducing the speed of the algorithm to increase the quality of solutions).

3. *Search space.* The type, number and resolution of your orchestra (and the corresponding knowledge of sound samples) will greatly influence the final result (also remember that some configurations are impossible to reconstruct, such as trying to reconstruct a low-pitched sound *pianissimo* sound only with high-pitched *forte* sounds).

4. If you have some doubts on your criteria selection, you can always give the *AutoPilot* a shot to try to optimize the criteria with the higher variability.

5. Finally, remember that different targets will require different settings, depending on the temporal nature, content and complexity. The same parameter settings will usually not be the optimal, even for two quite similar targets. So the best way to find good results is to explore the parameters of the system.

5.1.7 I added my own sounds but they are never used in orchestration

Several points can lead to the impossibility for the system to use your personal sounds (even though you added the folder in the explorer tab and the analysis went all the way through).

1. *Nomenclature.* For pitched sounds, you should first check if the sounds you added followed the required nomenclature. Indeed the system requires the sounds filename to be in the following form
   <Instrument acronym>-<playing style>-<note>-<dynamics>.
   For instance, adding samples of a *Bassoon* (with the acronym *Bn*) playing in *ordinario* should all be in the same folder and named in the following way
   Bn-ord-A4-mf.wav
   Bn-ord-C#2-ff.wav
   [...]

2. *Unpitched.* If you are trying to orchestrate with unpitched sounds or sound without any clear pitch or annotation (no note in the name), you *need to remove the "harmonic filters".* This checkbox can be found on the top left of the "*Solutions*" panel. This will ensure that the algorithm

3. *Orchestra.* Make sure that you added at least one instrument corresponding to your personal sounds folder.

4. *Filters.* Make sure that your filters allow to use your own sounds ! (That the corresponding *source, instrument* and *playing styles* are allowed). If you want to check if the algorithm is indeed able to use your sounds, *create a set of filters that allows only the use of these sounds* (by unchecking all other sources and playing styles).

5.2 Reporting a problem or sharing an idea

In order to facilitate the exchange of bug reports, exchange of idea and collaborative thinking, we encourage any kind of idea, report on errors, bad or good results, desired features, ideas, remarks and any other topic. Please report any unsolved problem, ideas or even cool sounds to*:*
http://forumnet.ircam.fr/user-groups/orchids

List of Figures