

International Music Information Retrieval Systems Evaluation Laboratory (IMIRSEL): Introducing D2K and M2K

{This document is a slightly revised version of the M2K demo handout provided at ISMIR 2004, Barcelona, in October 2004. Revised 13 December 2004.}

Overview:

The IMIRSEL (International Music Information Retrieval System Evaluation Laboratory) project provides an unprecedented platform for evaluating Music Information Retrieval (MIR) and Music Digital Library (MDL) techniques, by bringing together large corpora and significant computational resources with the necessary rights management and technical infrastructure to support a variety of MIR/MDL research areas. The standardized research collection being deployed represents a large and diverse corpus of musical examples, which we are hosting in our secure environment for use in evaluating MIR/MDL algorithms. The computational backbone and resource access of the IMIRSEL system is handled through the Data to Knowledge (D2K) machine learning and data mining environment developed by the Automated Learning Group (ALG) of the National Center for Supercomputing Applications (NCSA) at the University of Illinois Urbana-Champaign (UIUC).

Features of D2K:

D2K is a visual programming environment that allows for rapid prototyping and algorithm development. It is freely available on the web at:

<http://alg.ncsa.uiuc.edu/do/tools/d2k>


Some key advantages of the D2K environment include:

- A visual, modular programming environment to cut development time
- Support for distributed computing across multiple processors
- Capability of code reuse and sharing
- All components written in JAVA for maximum flexibility and portability
- Comes prepackaged with a large number of modules for machine learning and data mining specific to methods such as Bayesian classification and decision trees, to name a few

Potential users are encouraged to download to D2K and start using it locally on their own machines with their own datasets. Once the vast corpus of musical examples of the IMIRSEL project is made available to the community, users will be afforded the opportunity to take full advantage of the parallel computing capabilities of D2K, and run their itineraries on powerful computing clusters on datasets in the gigabyte, and even terabyte, scale.

Programming in D2K:

The development of algorithms using D2K entails wiring together computational modules into programs called *itineraries*, which represent dataflow between modules. These itineraries can then be run, or nested within other itineraries and used as modules, allowing for the development of itineraries with arbitrary complexity. As an example, the development of a real-time spectral analyzer with audio playback and visualization will be shown. Starting with a blank itinerary, modules can be selected from the module list to the left-hand side, and simply 'dragged and dropped' into the itinerary. Different module types are color coded with different visual icons. The core module types include: input modules (green), output modules (tan), data preparation modules (purple), compute modules (blue), visualization modules (brown), and user interface modules (red). The visual icons for the different module types are shown below. Moreover, to aid in module selection, the component info bar in the lower right corner can be expanded to show the documentation of modules as they are selected in the module selection area, or within the itinerary block diagram.

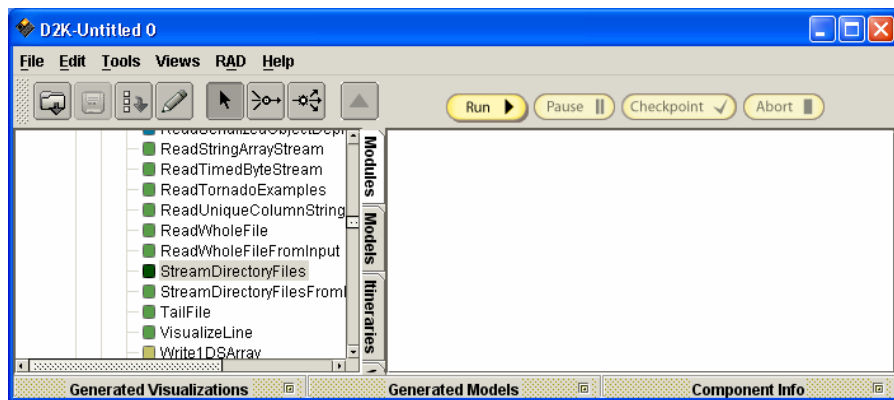


In this particular example, the required modules are:

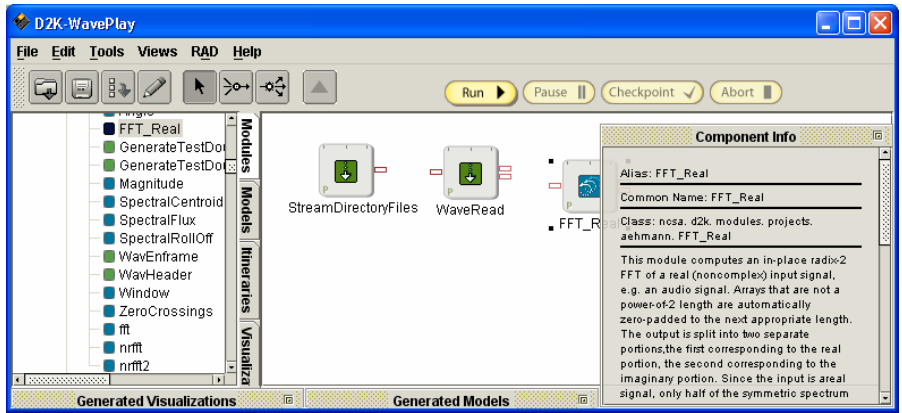
- *StreamDirectoryFiles*: Streams the names of all files in a given directory. Filters can be set, to only stream .wav files, for example
- *WaveRead*: Reads a .wav format file, and outputs the PCM stream as double precision floats between -1 and 1.
- *PlayDoubleAudio*: Plays the audio stream.
- *WavEnframe*: Segments the audio stream into short, overlapping frames.
- *Window*: Applies a user selectable window (e.g. Hamming) to the audio frames.
- *FFT_Real*: Performs a Fast Fourier Transform on the audio frames.
- *Magnitude*: Computes the magnitude of the spectrum.
- *VisualizeLine*: Creates a line plot of the spectrum.

Making the Spectrum Analyzer D2K Itinerary Example:

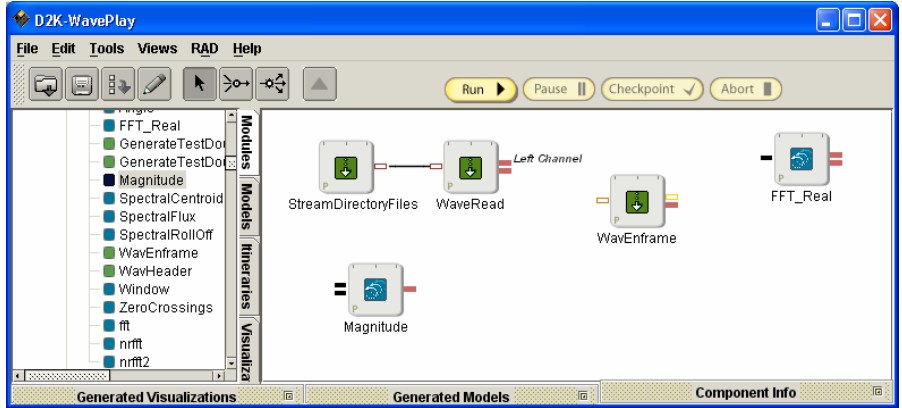
Shown below is a blank itinerary with module selection handled in the menu to the left.



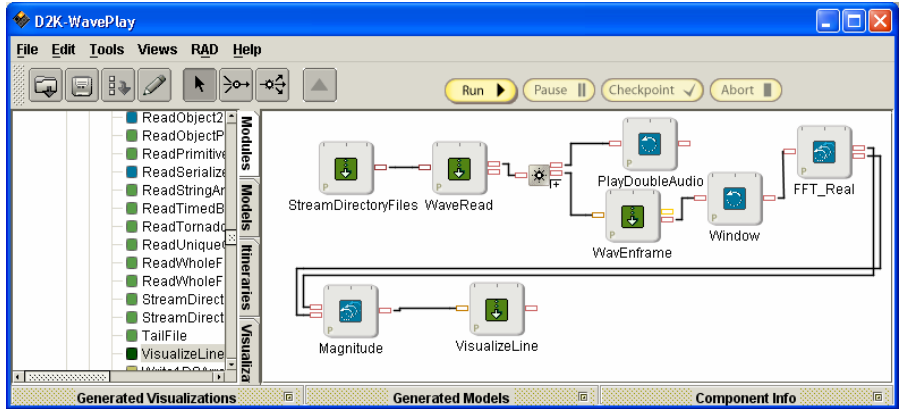
Modules are simply dragged and dropped into the itinerary. The component info panel is shown below, which provides information and documentation of the modules and their inputs and outputs.



Modules can then be wired together by simply putting the mouse over the proper tab and dragging a line to the desired connection point. As the mouse is held over an input/output connector, a label pops up as shown for the module *WaveRead*, whose first output is the left channel of the audio.




Shown below is the finished itinerary of the real time spectral analyzer with audio playback, completed in just five minutes!



Many modules have user definable parameters that can be set by double clicking on the desired module. For example, the *Window* module allows the user to select a window type to apply to an audio frame.



Once the itinerary has been constructed, and all pertinent parameters set, the program can be run by pressing the Run button at the top. Once the itinerary is started, the music is heard and a simple instantaneous time-varying spectrum shown for each frame.



Developing D2K Modules:

Individual modules to be used in D2K are coded in JAVA. Because of the standardized nature of the modules, a large percentage of the source code of existing modules can be reused in the creation of new ones. In general, labels and data types for the inputs and outputs must be set, along with necessary parameter fields. Within each module, exists a *doit()* method in which most of the module's computation takes place. This method is triggered by the presence of an input, after which the desired computation takes place with the results pushed to the output.

Future Work:

A large suite of modules and itineraries pertinent to music information retrieval and processing are currently being developed by the IMIRSEL team. These modules and itineraries will be made publicly available to the community. In addition, methods of interfacing D2K with other programming environments such as MATLAB, C/C++, and standalone applications are being explored, in the hopes of appealing to a large audience of MIR researchers.

Links:

<http://alg.ncsa.uiuc.edu/do/tools/d2k> - D2K download and information

<http://music-ir.org/evaluation> - Check here for links to M2K modules and itineraries developed by the IMIRSEL team and the MIR community.

Project Leaders:

Project Principal Investigator: J. Stephen Downie - *jdownie@uiuc.edu*
Key Developers: Andreas Ehmann
David Tcheng
Joe Futrelle
Kris West
Paul Lamere

Acknowledgements:

This work is supported by the Andrew W. Mellon Foundation and the National Science Foundation (IIS 0327371 and IDM 0328471). We also wish to thank the MIR/MDL community and our content-providers.