

## MAKING SOUND WITH NUMBERS, SIX YEARS LATER

*Nicola Bernardini, Damien Cirotteau, Free Ekanayaka and Andrea Glorioso*

The AGNULA Project  
<http://www.agnula.org>  
{nicb|damien|free|sama}  
@agnula.org

Media Innovation Unit  
Firenze Tecnologia  
Palazzo Borsa Merci  
Volta dei Mercanti, 1  
I-50122 Firenze  
Tel: +39335/706 3282  
Fax: +3906/233 200 750

<http://miu.firenzetecnologia.it>  
{nicb|damien|free|sama}@miu-ft.org

### ABSTRACT

In the first edition of the DAFx Conferences an extensive tutorial on professional and research software devoted to sound and music making was presented. The present paper attempts a revision of the concepts expressed in that tutorial, focusing particularly to the aspects related to research and innovation fostered by a strong paradigm shift that has happened in the mean time: that of Free Software development. Of course, this paradigm shift has also had its difficulties and harsh spots, requiring many extra efforts in order to overcome them. This paper will try to describe these as well as to outline the current state-of-the-art in the field.

### 1. INTRODUCTION

In the first edition of the DAFx Conferences, DAFx'98, one of the authors of this paper, along with Davide Rocchesso, presented an extensive tutorial on professional and research software devoted to sound and music making[1].

Six years have gone by since then, and a revision of that paper is indeed necessary: while most of the concepts expressed therein are still completely valid, a major change has occurred in the general software infrastructure that underlies most professional and research sound and music endeavors – namely, the expansion of Free Software to an extent which was quite difficult to predict at that time. True, that paper was indeed aware that “a good 70% of the software mentioned in this document belongs to the open-source public domain” (cf.[1, p.200]) and a re-edited version of it published in the Journal of New Music Research did specify further that “a good 50% of that software is actually Free Software” (cf.[2])<sup>1</sup>: certainly, Free Software was already very much alive and widespread in 1998; however, it was difficult to foresee that the following conditions

- professional-grade audio quality drivers

<sup>1</sup>Incidentally, this addition is a clear example of the language shift that happened in these last years giving unprecedented attention to the specific licensing aspects of software distributed over the Internet. Nowadays, a much more precise terminology has been developed, and care must be used when adopting it (cf. <http://www.gnu.org/philosophy/philosophy.html#TerminologyandDefinitions> ).

- unrivalled performance of the kernel infrastructure
- standard professional audio/music application development

would have (ever) been met by the Free Software communities. Quite unexpectedly, all these conditions (and more) have become a solid reality six years, completely changing the perspectives of sound and music-making today and posing new problems and challenges.

Effectively, the possibility of using a software infrastructure in sound research and production (including academic research in music, acoustics, audio signal processing and professional music production) completely based on Free Software opens up new and unexpected possibilities:

1. users may choose liberally to use consumer-grade or professional-grade applications according to their own needs, skills and capabilities; no artificial commercial barriers are set to separate “domestic users” from “professional” ones;
2. users are allowed (and as a matter of fact, even encouraged) to use these tools in a creative way, adapting them to their needs instead of accepting passively the features (or lack thereof) of each application; this is a particularly relevant feature in the sound and music domain, where creativity is often at stake in a dialectic relationship with the available tools;
3. users may easily mingle directly with developers on the many communication channels offered by the Free Software community at large (mailing lists, IRC channels, etc.), thus promoting a knowledgeable use of tools and applications;
4. users may contribute actively to steer development of tools according to their needs, thus accumulating knowledge and skills in a much faster process than what is usually achievable in non-free software environments;
5. bugs and problems are reported and disclosed to public inspection; as such, not only they get found and fixed with shorter delays, but they allow users to build a much more trusting attitude towards their tools (something that has gone completely lost in non-free software environments — now caught instead in a sort of fatalistic “reboot” attitude).

But, of course, this paradigm shift has also had its difficulties and harsh spots, requiring many extra efforts (i.e. other than writing software itself that is) in order to overcome them.

## 2. WHERE HAS ALL THE NON-FREE SOFTWARE GONE?

The Free Software paradigm shift is only one side of the story. In the previous DAFx paper [1] many non-free software packages devoted to sound processing and control were mentioned. In fact, in 1998 Free Software in audio was still a small niche in the complete picture. The situation depicted in Sec.1 may hint at the fact that the current Free Software ruling may have succeeded or will succeed in putting non-free software out of business.

This is strictly non-true. The first reason is that Free Software is not here to put software houses out of business: rather, it promotes fair and liberal competition among equally fitted individuals and institutions.

The second reason is that non-free software devoted to music and sound has been put out of business by its own practices. The music and sound domain chronicles of the past 4-5 years have been filled with gory stories of buyouts and subsequent disappearances of most software houses operating in this field. One prominent example was the Emagic (producer of the *Logic* sequencer) buyout by Apple and subsequent discontinuation of the software production line for the Windows production line in 2002<sup>2</sup>: the impact of the decision apparently affected several thousands Windows users, who discovered how Emagic would "cordially invite all Logic Windows users to join us [i.e. Emagic] on the Macintosh"<sup>3</sup>. Windows users were so upset by this decision that they went as far as signing a petition against it<sup>4</sup> to no avail whatsoever.

This is just one of the many examples available which cannot all be listed here for lack of space. For the same reason we will avoid the issue of all the software houses put out of business by impossible deals contracted with their reference operating system or hardware provider as well as that of the changing hands of software houses usually meaning vast changes in marketing targets, etc. At any rate, there have been so many cases now for non-free sound and music software brutally dumping their user communities that it is fair to assume that this has been the one of the most powerful motivations to push musicians and sound software users towards a conscious and aware use of Free Software. Certainly, making sound with numbers is possible today using Free Software exclusively. This paper strives to show why and how.

## 3. HISTORICAL BACKGROUND

As we already wrote, at around the same time than [1] was being conceived and written, the situation of sound/music Free Software applications had already reached what could be considered well beyond initial pioneering stage. A website, maintained by musician and GNU/Linux enthusiast Dave Phillips, was already collecting all possible sound and music software running on GNU/Linux

<sup>2</sup>cf. <http://www.beatmode.com/news/emagic/>

<sup>3</sup>This text appeared in a news originally on <http://www.emagic.de/english/news/>; of course, Emagic has deleted this embarrassing news a long time ago. However, the text in its integrity was pasted into the <http://www.beatmode.com/news/emagic/> page, so it is still available for evaluation.

<sup>4</sup>cf. <http://www.petitiononline.com/aewp/petition.html>

architectures<sup>5</sup>. That collection would already count hundreds of applications covering all possible sound/music fields and needs and particularly the broad categories already expressed in [1]: languages for sound processing, in-line sound processing, time domain graphical editing and processing, analysis/resynthesis packages, interactive graphic building environments, pedagogic software and processing libraries, plugins and toolkits.

At that time, the biggest problem was that all these applications were dispersed over the Internet: there was no common operational framework and each and every application was a case-study by itself. Free Software binary distributions were young: a quick glance over the Internet shows that full binary distributions of Free Software date back to the summer of 1993, Slackware being the first well-know one<sup>6</sup> in July 1993. Debian came second<sup>7</sup>, created in August 1993, while Red Hat was created one year later<sup>8</sup>. As a time-line reference, the GNU project startup declaration by Richard Stallman dates back to September 27, 1983<sup>9</sup> while Linux was announced for the first time by Linus Torvalds on October 5, 1991<sup>10</sup>. Thus, in 1998, while [1] was being written, no binary distribution was providing a solid coverage of multimedia applications and every sound/music fan/lover/enthusiast/researcher/professional had to get hold of Free Software sources, match dependency requirements (i.e. library versions, etc.), compile them, possibly debug/port them and then finally use them for whatever purpose.

A natural development followed [1] shortly after<sup>11</sup>, musician, composer and programmer Marco Trevisani proposed to a small group of friends (Nicola Bernardini, Maurizio De Cecco, Davide Rocchesso and Roberto Bresin) to create LAOS (the acronym of *Linux Audio Open Sourcing*), a binary distribution of all essential sound/music tools available at the time including website diffusion and support. LAOS came up too early, and it did not go very far.

In 2000, No Starch Press published a book by Dave Phillips devoted to sound/music in the GNU/Linux environment[3]. Besides providing excellent tutorial guiding, Phillips' book developed at greater length the taxonomy introduced by [1] concentrating mainly on Free Software applications: times were ripe for a successful attempt. And indeed, when Marco Trevisani proposed (this time to Nicola Bernardini, Günter Geiger, Dave Phillips and Maurizio De Cecco) to build DeMuDi (*Debian Multimedia Distribution*)<sup>12</sup>, an unofficial Debian-based binary distribution of sound/music Free Software.

Nicola Bernardini organized a workshop in Firenze, Italy at the beginning of June 2001, inviting an ever-growing group of supporters and contributors (including: Marco Trevisani, Günter Geiger, Dave Phillips, Paul Davis, François Déchelle, Georg Greve, Stanko Juzbasic, Giampiero Salvi, Maurizio Umberto Puxeddu and Gabriel Maldonado). That was the occasion to start the first concrete DeMuDi distribution, the venerable *0.0 alpha* which was

<sup>5</sup>The remnants of the website from those times are still visible in some old and forgotten mirror, such as <http://sunsite.univie.ac.at/Linux-soundapp/top.html>, while the current edition of it may be found at <http://linux-sound.org/top.html>

<sup>6</sup><http://www.slackware.com/announce/1.0.php>

<sup>7</sup><http://www.debian.org/intro/about#history>

<sup>8</sup><http://www.redhat.com/docs/manuals/linux/RHL-6.2-Manual/getting-started-guide/ch-history.html>

<sup>9</sup><http://www.gnu.org/gnu/initial-announcement.html>

<sup>10</sup><http://groups.google.com/groups?selm=1991Oct5.054106.4647%40klaava.Helsinki.FI>

<sup>11</sup>The earliest evidence being a personal mail dating back to November 1st 1998

<sup>12</sup>Personal mail exchanges on September 30th and October 23rd, 2000.

then quickly assembled by Günter Geiger with help from Marco Trevisani. A bootable CD-version was then burned just in time for the ICMC 2001 held in La Habana, Cuba, where Günter Geiger and Nicola Bernardini held a tutorial workshop showing features, uses and advantages of DeMuDi[4].

Practically at the same time<sup>13</sup>, Fernando Lopez-Lezcano was giving life to the PlanetCCRMA initiative[5]. PlanetCCRMA was initially conceived as an internal CCRMA service to promote the usage of GNU/Linux running on Red Hat boxes in computer music courses. Its success quickly surpassed the walls of Stanford University to be adopted by many users worldwide, thus providing a strong alternative (albeit very different in scope and intentions) to the DeMuDi distribution.

On November 26, 2001 the European Commission awarded the AGNULA (A GNU/Linux Audio distribution) Consortium (composed by the Centro Tempo Reale, IRCAM, the IUA-MTG at the Universitat Pompeu Fabra, the Free Software Foundation Europe, KTH and Red Hat France) with consistent funding for an accompanying measure lasting 24 months (IST-2001-34879). This accompanying measure, which was terminated on March 31st 2004, gave considerable thrust to the AGNULA/DeMuDi project providing scientific applications previously unreleased in binary form and a Red Hat based distribution parallel to the Debian (termed AGNULA/ReHMuDi). After the funded period, Media Innovation Unit, a component of Firenze Tecnologia (itself a technological agency of the Chamber of Commerce of Firenze) has decided to partly fund further AGNULA/DeMuDi developments. AGNULA has constituted a major step in the direction of creating a full-blown Free Software infrastructure devoted to audio, sound and music, but there's much more to it: it is the first example of a European-funded project to clearly specify the complete adherence of its results to the Free Software paradigm in the project contract, thus becoming an important precedent for similar projects in the future.

Nowadays, Free Software multimedia distributions are enjoying an outstanding success: new attempts of different size and purpose are started all over the world, and the future of Free Software for sound and music does indeed look bright. Among these, it is worthwhile to mention the [dyne:bolic] distribution<sup>14</sup>, another 100% Free Software distribution which enjoyed some European Community funding too through the sponsorship by the Public-VoiceXML project<sup>15</sup> (IST-2001-34546), and the Medialinux<sup>16</sup> distribution from the Open Source Lab of Virtual Reality & Multi Media Park in Torino, Italy.

#### 4. CURRENT STATUS

Given the successes of Free Software in the field of sound and music, we will attempt to give a short overview of the best Free Software available in each of the categories outlined in [1]. However, the reader must be aware that this is only a "best of" choice: distributions such as AGNULA/DeMuDi hold today a much larger variety of applications in each and every domain (cf. Figure 1).

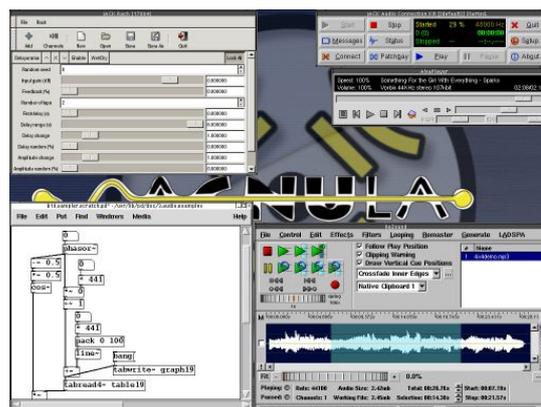


Figure 1: A typical AGNULA/DeMuDi desktop

**Sound Processing Languages:** the Sound Processing Languages domain provides an exemplary illustration of the paradigm shift explained in Sec.3. While standard customary applications did not change much in six years, it was rather the paradigm shift in licensing schemes to provide for real innovation here. Whether or not as a consequence of the creation of Free Software distributions for sound and music, some of the better known and used non-free applications of the past shifted to Free Software licenses. Two brilliant examples of this trend were *SuperCollider*, previously non-free and running on PowerPC platforms exclusively, which was re-licensed under the GNU GPL (GNU General Public License)<sup>17</sup> in September 2002, while the venerable *Csound*<sup>18</sup> music compiler was re-licensed under the GNU LGPL (GNU Lesser General Public License) around the month of May 2003.

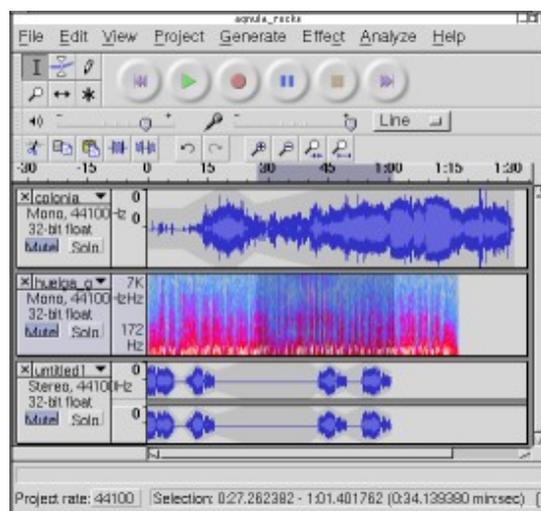


Figure 2: The Audacity editor

<sup>13</sup>According to its oldest changelog which dates back to September 14, 2001 (cf. <http://ccrma-www.stanford.edu/planetccrma/software/changelog2001.html>), that is the ending day of the ICMC 2001 Conference in La Habana.

<sup>14</sup><http://www.dynebolic.org>

<sup>15</sup><http://www.publicvoicexml.org>

<sup>16</sup><http://www.opensourcelab.it/article/articleview/7/1/9>

<sup>17</sup>An explained list of Free Software licenses may be found at <http://www.fsf.org/licenses/licenses.html#LicenseList>

<sup>18</sup><http://www.csounds.com>

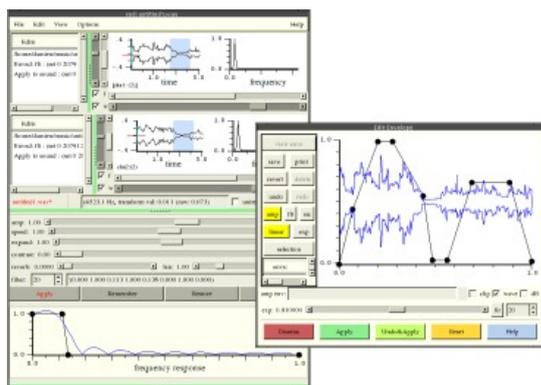


Figure 3: The Snd editor

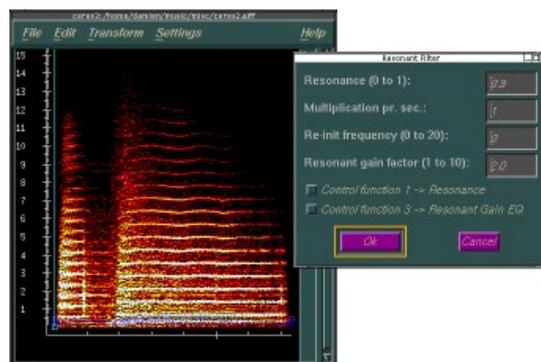


Figure 5: Ceres3, an Audiosculpt-like application



Figure 4: Ardour, the hard-disk recorder

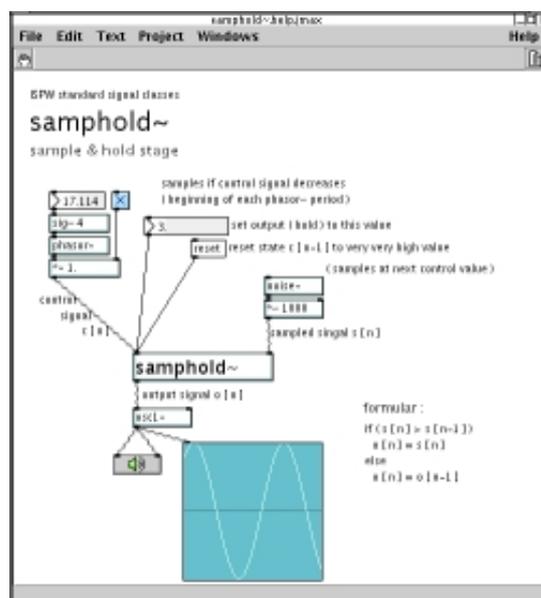


Figure 6: jMax

**Time-domain Graphical Editing and Processing:** Free Software graphical editing in the time domain currently sports among the best applications in the field: *Audacity* (cf. Figure 2), *Snd* (cf. Figure 3). A major enhancement was offered with the large scale multi-track Free Software hard-disk recorder *Ardour* (cf. Figure 4) which is directly competing with non-free software suites whose cost runs around the three-four figures.

**Analysis/Resynthesis Packages:** Most of the functionalities of *AudioSculpt*, IRCAM flagship non-free application described in [1] have been replicated into *Ceres3* (cf. Figure 5), a powerful Free Software editors in the spectral domain. Another powerful Analysis/Resynthesis package is *CLAM* [6], developed at Universitat Pompeu Fabra over the work previously done to invent and create the spectral modeling system *SMS*[7].

**Interactive Graphical Building Environments:** Free Software is very strong in this field with two well developed applications which have been enjoying a tremendous success for years: *jMax* (cf. Figure 6), direct incarnation of the *Max/MSP* environment, and *Pure Data* (better known as *PD*)[8]. Still, the creation of exclusive Free Software distributions such as *AGNULA/DeMuDi* has allowed to face and clarify licensing problems with both applications[9].

**Compositional/Expressive Tools** Other application packages that became Free Software in some recent past concerned Composition (*OpenMusic* — cf. Figure 7, developed at IRCAM for the MacOS platform and re-licensed into GNU GPL in July 2000) and Expression/Analysis Resynthesis (*Director Musices*, developed by the staff in force at KTH and recently licensed under the GNU GPL).

## 5. TECHNICAL DETAILS

All multimedia systems require very low-latency response in order to achieve real-time, and proper application precedence. While not particularly designed for real-time low-latency tasks, the Linux kernel, maintained by probably the largest development group in the world, made steep progresses in this field since 1999 onwards. A number of patches to the kernel were issued during this time, getting it to be a truly pre-emptive kernel<sup>19</sup> with very low laten-

<sup>19</sup>cf. <http://www.tech9.net/rml/linux>

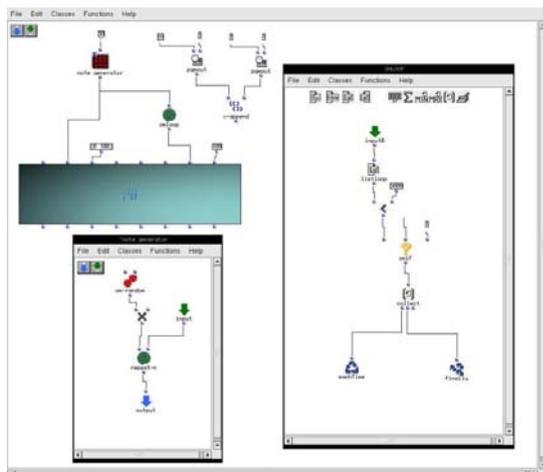


Figure 7: OpenMusic

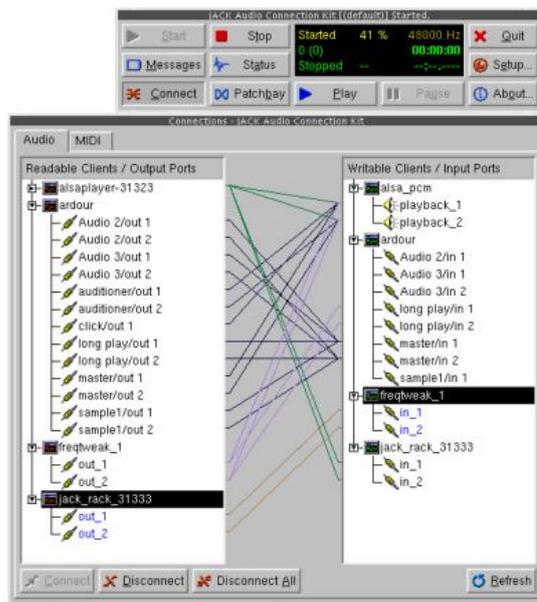


Figure 8: Jack Graphical Connection Client

cies<sup>20</sup>, achieving sub-2 msec latencies under most conditions<sup>21</sup>. Obviously, multimedia distributions currently offer patched kernels to achieve such latencies, effectively leading GNU/Linux to be the system of choice for multimedia endeavors.

But kernel latencies are only one of the innovative aspects of GNU/Linux system. Perhaps an even stronger aspect comes from the invention of *Jack*<sup>22</sup> by Paul Davis, early professional audio GNU/Linux supporter. *Jack* is a user-space low-latency audio server, written for POSIX conforming operating systems such as GNU/Linux. It can connect a number of different applications to an audio device, as well as allowing them to share audio between themselves. Its clients can run in their own processes (i.e. as normal applications), or can they can run within the JACK server (i.e. as a “plug-in”). *Jack* was designed from the ground up for professional audio work, and its design focuses on two key areas: synchronous execution of all clients, and low latency operation. Figure 8 shows a graphical connection client which hooks up to the *Jack* server. The development of *Jack* spurred a number of innovative applications such as the *Jack Rack*<sup>23</sup> (cf. Figure 9), essentially a host for *LADSPA*<sup>24</sup> plug-ins, which can then be applied to all audio paths within the system. Another brilliant application spurred from *Jack* is *Jamin*<sup>25</sup>, a module that was designed to perform high grade audio mastering of stereo input streams. Being a *Jack* client, it is easy to integrate it with the rest of the audio chain to allow the user to be able to change relevant parameters of the chain such as the mixing. As shown in Figure 10, the main components of this interface are a 1024 band hand drawn EQ, a 30 band graphic EQ, a 3 band peak compressor and a lookahead brickwall limiter. Furthermore, it allows smooth transition between different user’s presets to help the mastering process.

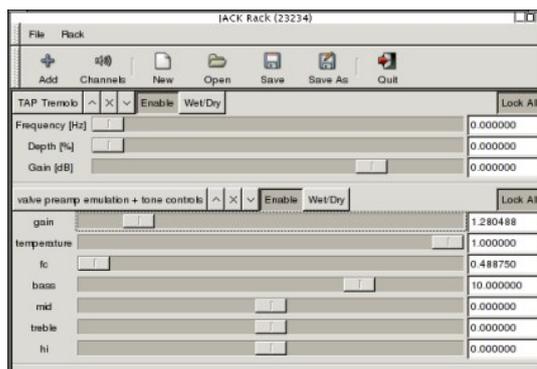


Figure 9: The Jack Rack client

## 6. CONCLUSIONS

We hope to have shown convincing evidence that the past six years have witnessed a paradigm shift (from non-free to Free Software for professional sound and audio applications) that will probably lead to many changes in habits and practices in sound research and music making. While many difficulties have been overcome by the amazing work of voluntary developers, it is clear that many more are still waiting to be solved. As a conclusive statement, we can try to list some of them:

1. there is still a lot of undergoing duplicate work: while projects such as AGNULA, PlanetCCRMA or MediaLinux are a good step towards a higher coordination of forces, many different programs basically performing similar tasks; this problem is actually double:
  - (a) developer’s resources get wasted;
  - (b) users must perform the additional task of selecting the proper application to suit their needs; this is often

<sup>20</sup>cf. <http://people.redhat.com/mingo/lowlatency-patches/>  
<sup>21</sup>cf. <http://www.gardena.net/benno/linux/audio/rtc2048-cpu80/2048.html>  
<sup>22</sup> <http://jackit.sourceforge.net/>  
<sup>23</sup><http://arb.bash.sh/rah/software/jack-rack/>  
<sup>24</sup>*LADSPA* stands for “Linux Audio Developer’s Simple Plug-in API”, cf. <http://www.ladspa.org/>  
<sup>25</sup>*Jamin* stands for “JACK Audio Connection Kit Audio Mastering interface”, cf. <http://jamin.sourceforge.net/>



Figure 10: Jamin

a hard and time-consuming job which musicians do not like to do;

2. most musicians tend to stick to their practised habits; thus, inertia to actually achieve a complete migration of most musicians is a much harder task than just building Free Software-based multimedia distributions;
3. in the research domain, a set of reference guidelines on the tools to be used and on the practices to follow still remain to be discussed and defined;
4. the professional music performance domain still lacks a common infrastructure that would lower the complexities of such projects; this infrastructure is missing in the non-free software realm too, where it is probably impossible to be achieved. Free Software constitutes a very good starting point to design and constitute this infrastructure – but this is something that is still to be done.

All in all, while unexpectedly outstanding leaps have been accomplished in the past six years, probably more achievements are ahead of us.

## 7. ACKNOWLEDGEMENTS

As the reader may expect, projects such as DeMuDi and AGNULA are the result of the common effort of a very large pool of motivated people. And indeed, giving credit to any deserving individual that contributed to these projects would probably fill completely the space allotted for this paper. Therefore, we decided to make an arbitrarily small selection of those without whose help DeMuDi and AGNULA would not probably exist. First of all, we would like to thank Richard Stallman, without whose effort Free Software would not exist at all; Linus Torvalds, who contributed the operating system we all got to depend on in the last dozen of years; and Ian Murdock who started Debian, the wonderful packaging infrastructure AGNULA/DeMuDi is based on. Then, Marco Trevisani, who has been pushing the envelope of a Free audio/music system for years, Dave Phillips, Günter Geiger, Fernando Lopez-Lezcano, François Déchelle and Davide Rocchesso: all these people have been working (and still work) on these concepts and ideas since the early days. Georg Greve, the president of Free Software Foundation Europe, has guided us through the difficulties of carrying out a rigorous Free Software project: may his

patience be awarded by the warmest thanks. Other people that deserve our gratitude are: Philippe Aigrain and Jean-François Junger, the European Commission officials that have been promoting the idea that AGNULA was a viable project against all odds inside the Commission itself; Dirk Van Rooy, later AGNULA Project Officer, Marc Leman and Xavier Perrot, patient AGNULA Project Reviewers; Luca Mantellassi and Giovanni Nebiolo, respectively President of Firenze’s Chamber of Commerce and CEO of Firenze Tecnologia, for their support: they have understood the innovative potential of Free Software much better than many so-called open-source evangelists. Finally, last but not least, Anna Meo and Nicola Giosmin, close collaborators for many years in all these (and other) endeavors, deserve all our endless esteem and gratitude: without them, a lot of these achievements would have made our life much harder than it has been.

## 8. REFERENCES

- [1] N. Bernardini and D. Rocchesso, “Making Sounds with Numbers: A Tutorial on Music Software dedicated to Digital Audio,” in *Proceedings of the 1st International Conference on Digital Audio Effects - DAFx98*, Barcelona, Spain, 1998, pp. 192–201, Institut de l’Audio Visual, Universitat Pompeu Fabra.
- [2] N. Bernardini and D. Rocchesso, “Making Sounds with Numbers: A Tutorial on Music Software dedicated to Digital Audio,” *Journal of New Music Research*, vol. 22, no. 5, 2002.
- [3] D. Phillips, *Linux Music & Sound*, No Starch press, San Francisco, USA, 2000.
- [4] François Déchelle, Günter Geiger, and Dave Phillips, “Demudi: The Debian Multimedia Distribution,” in *Proceedings of the 2001 International Computer Music Conference*, San Francisco USA, 2001, ICMA.
- [5] F. Lopez-Lezcano, “The PlanetCCRMA Software Collection,” in *Proceedings of the International Computer Music Conference 2002*, Goteborg, Sweden, 2002.
- [6] E. Robledo Arancio, R. Hinojosa, and M. de Boer, “A c++ development platform for real time audio processing and synthesis applications,” in *Proceedings of the 5th International Conference on Digital Audio Effects - DAFx02*, Hamburg, Germany, pp. 117–122, University of the Federal Armed Forces.
- [7] X. Serra and J.O. Smith, “Spectral modeling synthesis,” in *Proceedings of the International Computer Music Conference, 1989*, Columbus, Ohio, USA, 1989.
- [8] M. Puckette, “Pure Data,” in *Proceedings of the International Computer Music Conference, 1996*, San Francisco, CA. USA, 1996, pp. 269–272, International Computer Music Association.
- [9] The AGNULA staff, “1.4.1 - non/libre software recommendations,” Tech. Rep., The AGNULA Consortium, 2002.