The Daïs: A Haptically Enabled NIME for Controlling Physical Modeling Sound Synthesis Algorithms

Pelle J. Christensen
Aalborg University
Copenhagen, Denmark
pjch17@student.aau.dk

Dan Overholt
Aalborg University
Copenhagen, Denmark
dano@create.aau.dk

Stefania Serafin
Aalborg University
Copenhagen, Denmark
sts@create.aau.dk

ABSTRACT

In this paper we provide a detailed description of the development of a new interface for musical expression, the *daïs*, with focus on an iterative development process, control of physical models for sounds synthesis, and haptic feedback. The development process, consisting of three iterations, is covered along with a discussion of the tools and methods used. The sound synthesis algorithm for the daïs, a physical model of a bowed string, is covered and the mapping from the interface parameters to those of the synthesis algorithms is described in detail. Using a qualitative test, the affordances, advantages, and disadvantages of the chosen design, synthesis algorithm, and parameter mapping is highlighted. Lastly, the possibilities for future work are discussed with special focus on alternate sounds and mappings.

Author Keywords

NIME, Haptic feedback, Physical Modeling

CCS Concepts

•**Applied computing → Sound and music computing;** *Performing arts;*

1. INTRODUCTION

The *daïs*[1] is a novel interface for musical expression designed with multiple goals. The first goal was to develop an electronic instrument with multiple documented iterations. The second goal was to produce a repository of code, CAD models and build instructions, such that the build can be replicated. The third goal was to create an interface suitable for controlling *physical modeling* sound synthesis algorithms. To avoid a common pitfall, the daïs was intended to be easily reproduced by others by providing diagrams, assembly instructions, and software through GitHub[2].

A common way in commercial interfaces to interact with virtual instruments is through knobs, faders and simple push buttons (keys). These types of interactions are likely preferred because of the easy of mapping them to a scalar value, as well as because of affordability. However, knobs and buttons can be problematic in their expressiveness and playability, especially when controlling multiple parameters at the same time. Synthesis algorithms based on physical modeling often have many parameters, and for more advanced models we might need to control multiple parameters simultaneously in order to play the model at all. The daïs was therefore designed to provide intuitive control of multiple simultaneous parameters.

2. CONCEPT / RELATED WORK

The interface for the daïs incorporates a design that captures the position and orientation of a disk suspended by elastic string, as shown in Figure 1. The disk can be moved around using hands and the full 3D description can provide six simultaneous parameters — 3D rotation and position. This style of interface may be interesting for multiple reasons. First, one would be able to control up to six parameters at once, suitable for physical modeling algorithms. Second, a big disk affords big movements of the arm and wrist, which provides visual appeal forlive performances. Third, we have not seen an interface of precisely this type before, so it offered a new and unstudied way of interacting with sounds.

2.1 Haptic Feedback

The player-instrument system (preferably) features a tight feedback loop where the actions of the performer is guided by sensory cues received from the instrument. Though the auditory modality is dominant in this interaction, the visual and haptic channels are also important. It seems that the haptic feedback channel is often overlooked when designing electronic instruments. According to Cook this is one of the main reasons for the lack of intimacy with electronic instruments compared to acoustic ones [6].

Multiple studies have been performed on the importance of haptics in acoustical instruments. Askenfelt and Jansson showed that the vibration in stringed instruments are perceptible and that they might assist the player in intonating correctly, or in the case of the piano, help with timing[1]. In another study, Chafe showed that the sensation of the vibrations of the string helps cello players produce stable oscillations and that adding vibrotactile feedback to an electronic instrument can make it easier to use[4].

This research suggests that the addition of haptics to the daïs might enable more precise and expressive playing. Moreover, since the physical models are simulating vibrating systems, it would be interesting to actually feel those vibrations.

Haptic feedback has been applied widely to interfaces controlling physically based sound synthesis. An notable project is The Cordis System from ACROE for which several haptically enabled devices has been developed, one of

[1]daïs is a middle English word referring to a platform in a room where one would seat royalties or other dignified people.
[2]https://github.com/PelleJuul/dais

Licensed under a Creative Commons Attribution 4.0 International License (CC BY 4.0). Copyright remains with the author(s).

NIME'20, July 21-25, 2020, Royal Birmingham Conservatoire, Birmingham City University, Birmingham, United Kingdom.

which is their Gesture Transducer, which uses a motor and force sensor convert forces between the real and virtual world [3]. A related project employed the Haply development kit[3], an open source interface for haptic feedback and control [10].

For his PhiSEM project, Perry Cook constructed multiple haptically enabled percussion instruments using physically informed synthesis [6, 5]. A small qualitative test of an interface imitating a maraca showed that it provided an "uncanny feeling of connection between gesture, sound, and feel", which Cook attributes to the haptics and synthesis method [6]. Cook also comments on *passive haptics* in instrument design — the haptics that arise due to the mechanical design of the instrument, present in the daïs through the suspended disk, which provides the sensation of something pushing back at you.

2.2 Related Instruments

The Touché by Expressive E [4] features a sturdy base, upon which sits a stadium-shaped plate which can be moved around in a fashion similar to a joystick. The controller has no sounds on its own, but can be used to control either software synthesizers or effects trough MIDI, or to send CV voltages to analog synthesizers [8]. The main similarity between the Touché and the daïs is that they both allow sound to be manipulated by moving a plate. Another instrument similar to the daïs is the Sensel morph[5]. The main similarity there is that it offers an interface in the form of plate that the player touches with their hands.

3. DEVELOPMENT PROCESS

3.1 Tools and Methods

For the mechanical design of the interface we use OpenSCAD, an open source software solution for designing 3D objects[6]. Unlike most other CAD software packages, OpenSCAD uses a specially designed domain specific programming language for describing geometry, with a standard library and functionality for building complex objects. For the second iteration a custom PCB was designed, using KiCad — a free software suite for schematic and PCB design. For the processor we chose the Bela platform, a small, single board Linux computer (BeagleBone Black) with additional hardware for audio IO.

Before using the Bela, audio algorithms were developed in a desktop environment. For rapid prototyping and debugging of algorithms the first author developed the *Pretentious Audio Library* (PAL), which bundles mature open source audio and GUI libraries in an easy to use package.

3.2 First Iteration

The goal of the first iteration was to verify the effectiveness of our design idea. A picture of the first build, made out of cardboard and hot glue, can be seen in Figure 1. (a). Disk position sensing was first attempted with magnets and Hall-effect sensors. However, since the disk of the daïs can move somewhat freely, and because Hall-effect sensors are sensitive to angle as well as the distance, the measured value is ambiguous.

Vibrotactile feedback of the disk was achieved using two small vibration transducers: voice coils with an attached mass that can be adhered onto objects to induce vibrations.

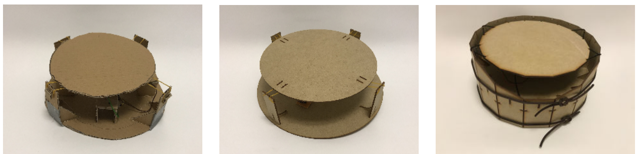

Figure 1: (a) 1st iteration (b) 2nd iteration (c) 3rd iteration

These were driven by the first channel of an external class-D stereo amplifier, the second channel was used to drive a small speaker for audio output. Testing this first iteration gave the impression that the suspended disk interface offered a new and interesting way of interacting with sound, because the interface afforded a kind of "wavey" movement of the arm. The vibrotactile feedback definitely made the disk more fun and intimate to interact with.

3.3 Second Iteration

The goal of our second iteration was to start working with mechanical design, and find an alternative to the Hall effect sensors. We decided to try optical proximity sensors, which emit infrared light and measure how much is reflected. The sensor used was a Broadcom APDS-9900 with 10 bits resolution and a range of 10cm. Like Hall effect sensors, it has an exponential response, which can cause issues with noise and resolution. This was mitigated by making sure that the disk was not too far from sensor, and applying a low pass filter to remove noise. The APDS-9900 has a maximal data rate of 2.5ms per measurement, which is plenty for control purposes. To capture the disk's movement and orientation in more detail, we detected: total depression (d) and rotation around the x and y (horizontal) axes (the normal). To calculate these it is necessary to detect three points on the disk's surface; to use three sensors.

A custom PCB was designed to ensure that the sensors were placed at known positions. A multiplexer was used to allow reading of all three sensors on the same I2C bus of a Bela. The sensors were calibrated to obtain measurements in millimeters by fitting sensor values to known distance values. Testing showed the computation of the d value was quite reliable. The computation of the normal was very sensitive to noise and errors in calibration, which made mapping of the rotation parameters nearly unusable.

We also experimented with mounting a piezo on the disk to capture more expressive interactions at audio rate, enabling the performer to use more nuance, e.g., scratching gestures on the surface. However, the vibrations of the disk were picked up by the piezo resulting in a feedback loop. We tried to solve this problem using a custom feedback suppression algorithm, but when that did not work optimally, we opted to instead move the piezo to the side of the instrument (an alternative interaction). Otherwise, interacting with this prototype was similar to the first iteration.

The synthesis algorithm (an early version of the final one) felt well suited for the interface and the vibrotactile transducers worked well with the stiffer fiberboard.

3.4 Third (final) Iteration

Since the solution with three proximity sensors was too noisy for the disk surface normal calculation, another sensor setup was developed. To measure d we now use a single proximity sensor, which eleminated the need for the multiplexer and the custom PCB, simplifying the design, software, and calibration routine.

For measuring the angle of the plate we use an Bosch Sensortec BNO055 IMU (intertial measurement unit), a sensor which includes the accelerometer/gyroscope sensor fusion

[3] https://www.robotshop.com/en/haply-development-kit.html
[4] https://www.expressivee.com/buy-touche
[5] https://sensel.com/pages/the-sensel-morph
[6] https://www.openscad.org/

algorithms internally. It can be configured with I2C commands, after which orientation and linear acceleration can be read at a rate of 100Hz. The chip is mounted to the disk with a piece of foam to isolate it from vibrations. The IMU provides orientation data (we use rotation around the x and y axis - the z axis currently unused).

We wanted this final iteration to look more finished. The aesthetic of the exterior affects how we perceive and interact with an instrument [9]. Therefore, electronic interfaces should also signal that they are instruments for creating art by having a beautiful, or at least nice, complete exterior. This meant having a nice way to mount all the sensors, route the wires, and a place to hide the Bela and amplifier board.

The new design is a fifteen-sided polygon (pentadecagon) as shown in Figure 1 (c). The construction is divided into two compartments: the lower compartment, which contains the electronic innards; and the upper compartment where the proximity sensor is mounted. Five of the wall pieces are lower than the others to make room for the wrist of the performer. The disk is mounted via elastic string, which isolates the vibrations of the disk from the piezo microphone.

4. AUDIO ALGORITHM

The sound synthesis algorithm used is a physical model of a bowed string. The method used is *finite difference schemes*, which is a way of numerically solving partial differential equations. For a through introduction to finite difference methods and the models and notation used here see [2]. The model we use is described by the differential equation

$$u_{tt} = c^2 u_{xx} - 2\sigma_0 u_t + 2\sigma_1 u_{txx} - F_b \phi(v_r), \qquad (1)$$

where $u = u(x, t)$ is the displacement of the string at point x and time t, c^2 is the wave speed determining the frequency of oscillation, σ_0 is general damping, and σ_1 is frequency dependent damping. The last term models the bow. F_b is the force at which the bow is pressed onto the string. The function $\phi(v_r)$ determines the frictional force applied to the string based on the relative velocity between the string and the bow. v_r is computed by

$$v_r = u_t - v_b, \qquad (2)$$

where v_b is the speed of the bow. ϕ can be defined in multiple ways, see e.g. [2, chapter 4.3], but we decided on

$$\phi(v_r) = \text{sign}(v_r)(\epsilon + (1-\epsilon)e^{-\alpha|v_r|}) \qquad (3)$$

as it was the most playable of the ones we tried.

To implement Equation (1) it is discretized as follows:

$$\delta_{tt} u_l^n = c^2 \delta_{xx} u_l^n - 2\sigma_0 \delta_{t\cdot} u_l^n + 2\sigma_1 \delta_{t-} \delta_{xx} u_l^n - F_b \phi(v_r), \qquad (4)$$

and Equation (2) as

$$v_r = \delta_{t-} u - v_b. \qquad (5)$$

5. PARAMETER MAPPING

F_b is mapped exponentially to the depression the disk of the daïs — a nice analogy to pressing a bow onto a string, especially combined with the haptics of the elastic string.

v_b is mapped to the tilt towards the player, which is maximally $a = 15\,\text{deg}$. For some playing techniques we need a mapping that allows $v_b = 0$ exactly, so when the disk is tilted away from the player, which means that $a \leq 1\,\text{deg}$, then we set $v_b = 0$. The mapping of v_b is unlike a real bow because the disk does not need to be moved continuously. To approach that, one map the angular velocity of the disk to v_b.

Figure 2: Test setup with the daïs, keyboard and speaker.

c^2 is mapped to the most recent pitch played on an external MIDI keyboard — a crude way of providing pitch control of our physical model, but very convenient to implement. Stability of the physical model is ensured by bounding the value of c^2. c^2 is also mapped to the side-to-side tilt of the disk to allow vibrato and pitch bends.

The pitch mapping causes an artificial, instant change in pitch which does not suit the organic sound of the string model. A better pitch mapping would be to introduce a fretting model that would introduce extra damping somewhere on the string such as in [11].

6. EVALUATION

To evaluate, a semi-structured interview was performed to see if test subjects would be able to understand the parameter mapping, what they thought of the expressiveness of the instrument, and to get feedback regarding the choice of sound and mapping as well as ideas for alternative sound designs and interactions.

6.1 Test Setup and Participants

Each interview started with one participant entering the room where the daïs was set up with an external speaker and a keyboard on the right as in Figure 2. Audio of the interview was recorded for future reference and notes regarding answers and interaction were taken in a spread sheet. The test then proceeded with the questions reported in the following section. The test was performed with six participants, all students within a music technology masters program. Each round took around ten to twenty minutes depending on how talkative the participant was.

6.2 Results and Discussion

What are your immediate thoughts when seeing this thing [the daïs]? / How do you think it is played? Four subjects expected to move the disk while controlling the pitch of the sound using the keyboard. One subject suggested a drum-like interaction and the last subject thought the tension of the elastic string might be a parameter. Two subjects found the instrument visually appealing.

What kind of sound do you think it will make? Two subjects expected a percussive sound. One expected a sustained sound. The other subjects expected digital sound. The subjects were also asked to do a verbal *sound sketch* [7]. Almost all subjects were reluctant to do this and no substantial observations came of their answers.

After playing the daïs without instruction: How did it feel *to play the instrument?* Two subjects played the instrument in a drum-like manner, while four other participants immediately pressed down the plate and were confused it did not make any sound right away. Subjects were mostly not able to figure out the mapping by themselves. One subject correctly identified the pitch bend mapping, and that pressing down and tilting towards/away would change the sound.

What kind of sound did it make? Most subjects described the sound as acoustic, string-like sound. One described the sound as bell-like. No participants directly suggested a violin or cello sound which might be due to the visual difference to traditional bowed string instruments.

After being instructed on how to play the daïs: How did it now feel to play the instrument? No subjects seemed to have trouble understanding the mapping. Three subjects immediately utilized the mapping and explored the sonic possibilities of the instrument. Two subjects were able to get some quite musical results — playing phrases and using the pitch bend. Some subjects mostly tapped the piezo mic to produce sounds. All subjects expected to be able to tap and then modulate the sound using the disk, but were disappointed because pressing down the disk kills the sound produced by tapping. One subject developed a technique of tapping on the side and then pressing the disk to produce a sustained tone with a strong attack, which is a technique we had not thought of previously. Most subjects described the playing experience as fun, enjoyable and intuitive once they got a hang of the mapping.

Can you think of any other sound or interaction style that would suit the interface? Many subjects suggested that a percussive sound could be suitable and that moving the disk around should modulate the sound. One suggested that a wind instrument sound could also be nice to try.

Do you have any other comments on the instrument or the testing procedure? One subject would like to see the daïs and the keyboard being built together into one object or maybe with continuous pitch control instead. It seems that multiple of the subjects were afraid to break something since most of them refrained from pressing the disk down very far. One subject commented that they liked the big, performative movements afforded by the interface.

Only one subject commented on the haptic feedback and stated that they really liked it. When inquired about the haptic feedback all subjects stated that they noticed it but did not actively think about it and that they really liked it. The fact that the subjects did not comment on the vibrotactile feedback by themselves can be interpreted in two ways: either the vibrotactile feedback does not affect the playing experience enough to merit a comment, or, the haptic feedback enriches the experience in such a natural way that you don't even think about it. We are inclined to think the second interpretation is the right because the interaction is similar to an acoustical instrument where vibrations of the instrument body is a natural part of the combined experience of playing the instrument.

6.3 Summary of Evaluation

The chosen mapping scheme is not at all intuitive and new users would have a hard time figuring it out without instruction. This is because the mapping was not designed to be intuitive, but to afford expressive playing once you know how it works. The short amount of time the test lasted for was not enough to learn the mapping properly. The exterior of the daïs suggests a percussive sound to some and an abstract synthy sound to others. The actual bowed string sound was not suggested by any test participants. The visual similarity to a drum has an impact on what users expect of the interface; a mapping to a percussive sound would be a natural choice for a new mapping. Once introduced to the mapping users found the instrument fun, enjoyable and natural to play. Because of the mapping and the physical modeling algorithm, new playing techniques and timbres can be discovered, and the instrument gets a life of its own, sometimes producing unexpected sounds. Finally, the haptic feedback adds another natural dimension to the playing experience and made the playing more fun and expressive.

7. CONCLUSION

Since some subjects were afraid to damage the interface, more time should be spent on the design to make the daïs more sturdy. In an extended evaluation, one could instruct participants in how much abuse the daïs is able to take, since it is more robust than it may appear. In the future, it will be interesting to test different sounds and mappings. A takeaway from the test is that a mapping to a percussive sound would be good, using the contact microphone as a trigger and the disk to manipulate the sound. One colleague suggested that granular synthesis would be a good choice since the disk could be used to navigate though a sound space in an intuitive way.

8. REFERENCES

[1] A. Askenfelt and E. V. Jansson. On Vibration Sensation and Finger Touch in Stringed Instrument Playing. *Music Perception: An Interdisciplinary Journal*, 9(3):311–349, Apr. 1992.

[2] S. Bilbao. *Numerical Sound Synthesis*. John Wiley & Sons, Ltd, Chichester, UK, Oct. 2009.

[3] C. Cadoz, A. Luciani, J. Florens, C. Roads, and F. Chadabe. Responsive Input Devices and Sound Synthesis by Stimulation of Instrumental Mechanisms: The Cordis System. *Computer Music Journal*, 8(3):60, 1984.

[4] C. Chafe. Tactile Audio Feedback. In *Proc. Intl. Computer Music Conf.*, Tokyo, 1993.

[5] P. R. Cook. Physically Informed Sonic Modeling (PhISM): Percussive Synthesis. In *Proc. of The International Computer Music Conference*, 1996.

[6] P. R. Cook. Remutualizing the Musical Instrument: Co-Design of Synthesis Algorithms and Controllers. *Journal of New Music Research*, 33(3):315–320, Sept. 2004.

[7] D. Rocchesso, G. Lemaitre, P. Susini, S. Ternström, and P. Boussard. Sketching sound with voice and gesture. *ACM Interactions*, 22(1):38–41, Jan. 2015.

[8] E. Simon. Improved haptic controller, eu patent ep3129981b1, 2018.

[9] L. Turchet. Smart Musical Instruments: Vision, Design Principles, and Future Directions. *IEEE Access*, 7:8944–8963, 2019.

[10] J. Villeneuve and J. Leonard. Mass-interaction physical models for sound and multi-sensory creation : Starting anew. In *Proceedings of the 16th Sound and Music Computing Conference*, pages 187 – 194, Malage, 2019.

[11] S. Willemsen, N. Andersson, S. Serafin, and S. Bilbao. Real-time control of large-scale modular physical models using the sensel morph. In *Proceedings of the 16th Sound and Music Computing Conference*, pages 151 – 158, Malaga, May 2019.