# Chapter 2

# Psychophysics of the Piano

> The moral for piano teachers is that so far as single notes are concerned, it does not matter how the pupil strikes the key, so long as he strikes it with the requisite degree of force. If this is right, the tone quality will be the same whether he strikes it with his fingers or even the end of his umbrella. As far as the scientist can see, that is all there is to the much debated problem on the piano touch.

–Sir James Jeans, from a lecture read to the English Piano Teachers' Association, Jan. 8, 1939 [52]

## 2.1 The piano: a paradox

The piano, when regarded both from the standpoint of an engineer and that of a musician, presents something of a paradox. The engineer, whose primary interests are learning from or improving upon the piano's design, will inevitably find it difficult to reconcile his convictions about this instrument with those of the musician, whose interests are centered around musical expression. The engineer points to the simple principles by which the piano produces sound and the correspondingly small set of controls over these principles which the piano makes available to its players at the keyboard. He underlines the fact that the piano is fundamentally a percussion instrument. The musician, on the other hand, points to the rich music which the piano can produce, nuanced not only in harmony and phrasing, but also in loudness and tone color. The expert pianist can even demonstrate independent control over each of these parameters. But given the degree of decoupling which the piano imposes between the pianist's fingers and its sound-producing mechanics, even the musician must concede that it is by no means apparent how one may use this instrument to produce evocative music.

In this introductory section, I will attempt to unravel the discussion between engineer and musician over the piano and the music it produces. I am interested in laying open this apparent paradox in order to motivate and enable the design of electronic musical instruments which more fully approach the piano in their expressive capabilities. If rapport could be established between engineer and pianist, the design of electronic instruments could in fact become a collaborative effort between the two, an activity which, arguably, has not yet occurred.

As already intimated, the investigations of this thesis center around a proposed re-design of electronic instruments, specifically, synthesizer keyboards. We envision a keyboard which offers the advantages of programmability, computerization, and mechanization, yet still realizes the full expressive capability of the grand piano. Not until we can re-engineer and modernize the piano without losing any of its capacity as an expressive instrument will we have invented a worthy successor in the form of a digital instrument.

A redesign of electronic pianos will naturally require that numerous design decisions be made; these decisions will in turn require quantitative bases. Thus, our analytical tools will be those of the engineer. Our endeavor, however, is completely accountable to and driven by the musician. Indeed, the extent to which this new instrument facilitates artistic expression for the musician will be the ultimate gauge of the success of a re-engineering of the piano. Our particular concern is the 'feel' of the piano at the keys, but before we address the topic of mechanical interactions between the pianist and piano, let us further motivate our redesign effort by revisiting the juxtaposition of views, the musician's and the engineer's.

### 2.1.1 Subdominant: The engineer's viewpoint

Compared to other musical instruments, the piano provides the musician with a very restricted set of controls over the parameters of sound. Unlike wind and bowed string instruments, which depend upon continuous excitation of a resonating body for the production of a tone, the piano offers no provisions for varying timbre [1] independent of intensity. For each pitch on the piano, there exists a one-to-one correspondence between intensity and timbre. The two are jointly determined by a key stroke. This fact has to do with the piano's percussive nature and will be further discussed below.

Unlike the voice (but similar to most other musical instruments), the piano is not capable of modulating the filtering properties of its resonating body. After the hammer-string interaction period is over, the sound of the piano evolves according to the static filtering properties of the string and soundboard.

---

[1] Timbre is a rather difficult term to define because it is a perceived quantity and further a multi-dimensioned quantity. The Acoustical Society of America has settled on defining timbre as *all tone qualities not already defined as pitch*. See [58]. I use the words timbre and tone color somewhat interchangeably to refer to the frequency spectrum of a tone.

The piano, along with its keyboard instrument relatives, does have one important advantage over other instruments: readily produced polyphony. It's immense reputation as an expressive instrument, however, cannot be fully accounted for by this one factor. The piano also enjoys status as an instrument of rich tone color. Its music is widely recognized as lyric, with a color palette practically on par with any string or wind instrument. Piano performances are characterized not only according to their range in intensity (loudness) but also according to their range in shading and timbre. In light of the percussive nature of the instrument, however, it really does come as a surprise that the piano should be revered for its ability to produce nuance in tone color. The phenomenon is especially astonishing since a pianist also lacks control over the precise pitches produced by the piano; these are under the control of the piano tuner. I will further illustrate the basis for the engineer's claim that the pianist has very limited control over the sound of the piano by considering the piano action as a kind of transform or map.

### Mapping from gesture to sound parameters

Instigated by key-presses, sounds are produced on the piano by percussive hammer strikes on strings. The period of interaction between hammer and string, during which a waveform is set up on the string, is very brief (about 2 milliseconds) [3]. This waveform quickly evolves into a standing wave which gradually leaks its energy into the soundboard and from there into the surrounding air. The physics governing the hammer-string interaction are complex, owing in large part to the compression-hardening properties of the felt covering the hammer. For a look into the ongoing research of the hammer-string interaction, see [40] and [13]. Our particular interest, however, lies in the fact that the key and the pianist's finger are completely decoupled from the hammer during its brief period of interaction with the string. The hammer flies free of the jack which initially propels it some 2.5 milliseconds before striking the string [3]. The pianist has no means at his disposal for controlling the tone or the evolution of tone after the hammer has left the jack, except through the damper. Thus all parameters of the tone must be set up by the pianist before tone onset. [2]

The piano action, then, can be regarded as a system which maps a keypress, initiated at a

---

[2] The possibility of the hammer supporting more complex motions (such as vibration modes in the hammer shank) has been refuted by the excellent experimental studies of Anders Askenfelt and Jens Jannson [2]. Vibration modes in the hammer shank were hypothesized as a means for independent control over intensity and timbre. If the motion of the piano hammer in flight could support more than just the rigid body mode, then perhaps the hammer-string interaction would give rise to various waveforms on the string as a function of the various relative mode energies, which in turn are presumably a function of the manner in which the key is struck. Thus a more impulsive key strike would introduce more energy in higher hammer vibration modes, and the hammer-string interaction would result in a more complex wave motion in the string and thus a brighter tone. The experimental results of Askenfelt and Jannson, however, indicate that the amplitude of the vibration modes of the modern hammer are so small as to be immeasurable and in any case have no effect on the tone. Interestingly, vibrations were in fact detectable on the more slender hammer shank of the historical forte-piano hammer.

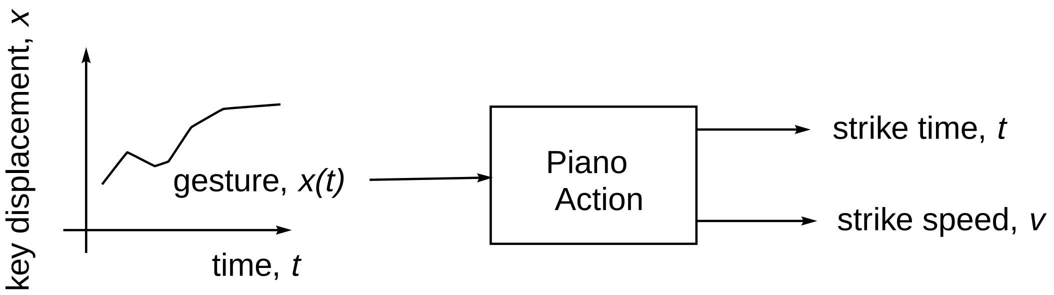

Figure 2.1: *Mapping from Gesture to Sound Parameters as Carried Out by the Piano Action*

certain instant, into a final hammer angular speed occuring at another certain instant. The keypress requires a graph for its full description. It is a function defined over a time interval. But rather than attempting to define it with either a force or a motion history, I shall simply refer to the keypress as a *gesture* since the human finger acts neither as a perfect velocity source nor a perfect force source. Figure 2.1 illustrates the mapping from an input gesture to two output scalars, the hammer strike time and hammer strike speed. Note that the hammer strike time is not directly tied to the initiation time of the gesture; the time interval between initiation and strike is strongly dependent on the shape of the graph of the gesture.

Beyond selecting a pitch by choosing a particular key, the pianist has only to select the two output parameters to completely determine a tone. With the selection of the speed, both the tone intensity and its timbre are jointly determined, for they are coupled, as mentioned above. Although a gesture at the key may be complex and may encode the complex expressive intentions of the performer, its action in the end is completely encapsulated by the two simple scalars to which it gives rise. The mapping from gesture into two scalars is itself complex insofar that it depends on the dynamics of the piano action, a rather daunting (on the surface) dynamical system which will be described in detail and modeled in Chapter 3. For example, similarly scaled gestures will not necessarily give rise to similarly scaled final hammer speeds. But the fact remains that for all practical purposes, the piano action plays a role which is a *severe data reduction.*

Finally, note that this mapping is many to one, that is, there are many ways to strike the key yet produce the same final hammer speed and strike time. The mapping is not, however, one to many. A certain gesture at the key produces a very repeatable final hammer speed and strike time. That the mapping performed by the piano action is repeatable and has as its output an event parameterized by two scalars gives credence to the engineer's claim that the operation of the piano action is 'simple'.

### 2.1.2 Dominant: The musician's viewpoint

The existence of a complex and highly developed piano technique among pianists, regardless of musical persuasions, seems to contradict the engineer's statement, that all that is involved is the selection of two scalars (strike time and strike speed) for each tone. One must wonder why there exist so many ways to hit the key and, especially, why certain techniques are recommended for their effectiveness at producing certain tone colors. Apparently the behavior of this instrument in the hands of its users is responsive enough to warrant a large amount of attention to skill development. Certainly a large part of that skill lies in selecting the two scalars, hammer strike time and hammer strike speed, but here I am speaking of responsiveness of the piano to the various techniques which are chosen for their coloristic effect, yet intensity invariance. The piano seems to reward its users with tones which vary in color, yet remain constant in intensity according to the use of these chosen techniques. Certain strategies adopted at the keyboard have results which lead the listener to believe that timbre and intensity are being varied independently.

The language which piano teachers speak when trying to aid their students in producing music at the piano is quite obviously built on the assumption that a pianist has independent control over intensity and timbre. Certain techniques are taught or encouraged with the explicit intention of making a passage soft yet dark, or loud yet airy. Piano teachers use such terms despite what they may know of the limitations of the mechanics of the piano action as outlined in the previous subsection.

The existence today of synthesizers and digital pianos allows us to make some further observations in support of the musician's viewpoint. The fact that a digital piano does not seem to reward a pianist's various techniques with independently varying sound parameters suggests that there is indeed something about the piano that has been missed in its representation in digital pianos. I will, however, delay discussions about modern electronic instruments until section 2.2.

### 2.1.3 Tonic: resolution

How then do we reconcile the positions of the engineer and pianist, if up until now our discussions have given them both support for their disparate claims? The engineer assures us that one scalar (hammer angular speed) cannot independently determine two effects (timbre and intensity). Intensity and timbre are coupled for each tone. The musician, on the other hand, claims to have independent control over intensity and timbre and demonstrates that control, even for the ears of the engineer!

To suggest that both the engineer and musician are both right would seem counterproductive, but that is essentially what I am going to argue here. The musician is merely right (making a truthful statement) about the issue as it occurs and presents itself in another domain: the perceptual domain. The musician is speaking of a perceived phenomenon rather than a physical phenomenon. It is thus with reference to a tenet of cognitive psychology that we resolve the conflict between the engineer and musician, and legitimize the paradox of the piano rather than deny it.

The proposition that psychophysical quantities may have rather remote relationships to the actual physical quantities which underlie them is not a new idea. Cognitive psychologists have defined a distinction between the physical world and the world as it is perceived. The perceived world is established by an individual based on incoming data filtered through the senses. These data are subject to transformations and processing by perceptual operators in the brain or the sensors themselves. The perceived world, however, exists in a manner which is no less legitimate than the real world. It can be operated upon, practiced with, extended, and explored. Furthermore, there exists a tendency by all humans to forget the distinction between the physical and perceived worlds, to accept the perceived world and in fact identify it as the physical world. This tendency to project the perceived world onto the physical world has been called "distal attribution" by [104].

The musician is speaking of this perceived world when he says that a piano tone's intensity and timbre can be independently varied. The fact that the musician is unable to distinguish between perceived reality and physical reality is due to distal attribution. Of interest to us here is the manner in which the transformation from physical parameters to perceived parameters occurs. If we understood this transformation, we could further exploit it in the design of electronic instruments. What parameters of piano music inspire its listeners to identify tones of the same loudness as having different tone color? We would like to identify the physically measurable parameters which the musician (or listener in general) identifies as timbre differences and attributes to certain qualities about key strikes.

I propose that what the musician hears as control over timbre is effected by careful control over timing. For a particular tone, the pianist is indeed able to select only two parameters beyond pitch: the hammer strike speed and the hammer strike time (although the time of tone damping is also important). I propose that by selecting these two parameters, a pianist can produce a percept in his or her listener of timbre being controlled independently of intensity. The contradiction between viewpoints is resolved if we assume that what the listener hears as timbre control is due to fine timing control in the physical domain. The percept depends to large degree on the consideration of not just a single tone, but a group of tones arranged in a musical phrase. By carefully governing the timing overlap of notes as they follow one another, the musician can evoke a certain percept in listeners which will not be labeled in terms of timing at all; it will be labeled by the ear and auditory perceptual centers in terms of the frequency domain, that is, timbre. Although it is not, in physical terms, independent control of timbre and intensity, it is perceived (and labeled) as such.

For example, an arpeggio played with disconnected notes may sound different in timbre than one played with slurred overlapping tones. Tone color differences do indeed exist in the perceived world, despite the fact that they are not supported by the physical phenomena to which they refer.

Finally, we remark that the musician can indeed make valuable contributions to the design of musical instruments, despite the existence of paradoxes between the physics and the perceived music of that instrument. The assumptions which develop in musical pedagogy deserve careful consideration from the engineer interested in instrument building. The languages spoken by the engineer and musician may both be acknowledged as truthful if we allow them to pertain to different domains, and harken to the relationships between those domains with psychophysical studies.

In the present work, I will not contribute directly to the discussion between engineer and musician. I leave these interesting psychophysical studies to future work and to the work of others. However, the premise that control over timing is recognized as control over timbre motivates a deeper consideration of the mapping which the piano action performs on the input gesture. If timing is so critical, it is worth considering issues such as sensitivity of that mapping to slight variations in the input gesture. The gradients of the mapping deserve careful attention in the design of an instrument. Furthermore, the energetic mechanical interactions between pianist and piano become very important to consider. Power exchanges between finger and key are intimately associated with timing. It also becomes plausible that certain techniques are superior for their robustness in the face of disturbances. Disturbances may arise, for example, from the lack of precise control or repeatable control over muscles. This discussion will be revisited in Section 2.5 below.

### 2.1.4 Literature Review: Paradox of the Piano

Although the topic of the paradox of the piano will not be investigated *per se* in this work, I have collected various sources in the literature which address this topic. The supposition that perceived timbre effects on the piano are due to timing and intensity variations is not new. The ballistic nature of the piano action is quite accessible and has long been understood in qualitative, if not in quantitative terms. I have not found, however, a significant body of literature which addresses the genesis of the timbre variation percepts. This introductory section may be considered a call to arms for this research. It is also hoped that an electronic instrument, as proposed and investigated in this work, could become a tool for such a study.

Many musical acoustics researchers have found the piano paradoxical enough to warrant bringing the issue into the lab for scientific investigation. In 1925, Otto Ortmann, of the Psychological Laboratory of the Peabody Conservatory of Music, published his book Physical Basis of Touch and Tone: an Experimental I
Ortmann conducted a very thorough investigation which included recordings of the key motion made

with traces left by a vibrating pitch-fork in smoked glass attached to a key. Ortman concluded: “The only factor directly influencing, or responsible for, the vibration of the piano string is the speed with which the hammer leaves the escapement.” [78].

In 1934, Harry C. Hart, Melville W. Fuller, and Walter S. Lusby, of the Electrical Engineering department of the University of Pennsylvania, encouraged by Professor Charles N. Weyl, and with the cooperation of a well-known concert pianist, Abram Chasins, conducted some experiments on the piano to determine if the tone as played by a mechanical key striker could be distinguished from that of a live pianist. Using human subjects as auditors, and camera recordings of string motions, they determined that indeed, single notes struck by a finger could be duplicated in every way by a mechanical striker [42].

Carl E. Seashore contributed further to the work of Ortman in 1939 [91]. Together with Tiffin, Seashore used an apparatus for recording hammer motions which they dubbed the Iowa Piano Camera. Seashore is to be commended for his keen understanding of the perceptual factors at play.

Sir James Jeans, aware of the work of Hart et al., sparked an uproar in the piano community with his reading of a lecture to the English Music Teachers’ Association on January 7, 1939, which included the quote about the umbrella handle which I used to open Chapter 2. Jeans’ lecture was reprinted in full on January 8th and 9th in the New York Times, and was followed up with an interesting article in the Times “This Week in Science” by W. Kaempfert on January 15th of 1939 [53]. Kaempfert, in an attempt to calm the apparent commotion and clear up misunderstandings, pointed out that the scientific results pertained only to a single strike: “No time should be wasted in achieving beauty in single notes, though a good deal of time may be profitably spent in acquiring technique as a whole.”

In 1950, J. Helmann, a piano pedagogue at the St. Petersburg Conservatory, took it upon himself to directly refute the claims of James Jeans and company in his book The Consciously Controlled Piano Tone [45]. Helmann bases his evidence on informal psychoacoustic experiments with his students as stimulus producers. Helmann documents an extensive technique of hand motions and speed/force control and a theory to back it up. Helmann was not clear, however, on the distinction between perceptual reality and physical reality.

The commercially very successful reproducing pianos, first introduced around 1910, further fueled the public debate over the mystic qualities of human touch. Reproducing pianos were a technical improvement on the player piano. In addition to the hammer strike times, the reproducing piano was able to record the hammer strike speeds with an extra two tracks on the piano roll. Likewise, the reproducing piano was capable of modulating the pneumatic actuation of the keys. This technology successfully competed with sound recording for about a quarter of a century. Many famous artists made recordings with these pianos. Surviving roles have been restored and used with refurbished reproducing pianos to make audio recordings with modern digital recording technology. These recordings are occasionally hailed today as very accurate reproductions original performances. But in fact, these recordings were usually edited to sound better than the originals, sometimes without the artist's involvement.

Today there exists another generation of reproducing pianos, these implemented with electronics and electromagnetic actuators. The Boesendorfer Recording 190 SE and the Yamaha Disklavier are examples. In contrast to the electronic instruments which will be discussed in the upcoming section, all of these instruments contain a piano action, enabling the recording of the hammer motion. Subsequent perfect reproduction is then possible through design, by trial and error, of an appropriate actuated 'gesture' at the key. So long as the strike speed and strike times of the hammer can be encoded and reproduced, the magic of the piano can be laid bare, as is clear from the above discussion. It is very important to note, however, that the original production of these events by a human player is made possible only by the existence of a piano action in the instrument. During reproduction (recording play-back), the actuators need not use the same gestures as the player, only some known (pre-determined) gestures which will produce the same hammer strike time and speed. The mapping from gesture to hammer strike speed and strike time is many-to-one.

## 2.2 Synthesizer Shortcomings

So called 'digital pianos' exist *en masse* on the market today. They cater to a significant portion of the keyboard-instrument market formerly occupied solely by the acoustic piano. Digital pianos are essentially synthesizers which feature the best available synthetic (usually sampled) piano sounds and a 'weighted' keyboard. Weighted keyboard is a simple keyboard action whose feel at the keys is made to roughly approximate that of the piano with the incorporation of springs, dashpots, or weights. Digital pianos offer various advantages over their acoustic predecessors: portability, sound-programmability, and occasionally, price.

In this author's opinion, digital pianos represent only a first cut at an electronic instrument modeled after the acoustic piano. Certainly the recent advent of affordable digital sound synthesis has made emulation of the acoustic response of the piano possible [3], but the touch-response of even the best weighted keyboards does not yet compare favorably to the touch-response of a grand

---

[3] Improvements in the computational reproduction of the piano sound are continually being made by numerous designers and computer music researchers. Most notably, a new synthesis technique known as ***physical modeling***, with several proponents at Stanford's Center for Computer Research in Music and Acoustics (CCRMA), [92] [100] promises great gains in the fidelity of audio emulation. The effects of string coupling through the bridge and soundboard and the effects of the nonlinear properties of the hammer felt on the sound can be elegantly integrated into such sound synthesis algorithms.

piano. A second shortcoming of digital pianos, intertwined in origin with the unsatisfactory touch-response, is their implementation of an overly simplistic mapping from gesture to produced sound. Both of these deficiencies factor into why these products cannot answer to the desires of classically (or otherwise highly) trained pianists. In particular, the limited success at emulation exhibited by digital pianos to date motivates careful attention into the dynamical behavior of the target system (the grand piano action) both in terms of its touch-response and the mapping which it performs from gesture to sound parameters. The work of this thesis is primarily intended to remedy the shortcomings of digital pianos in the area of touch-response, yet, as a byproduct of our approach, the improper mapping from gesture to sound parameters will also be amended. Section 2.2.2 will address in detail the mapping from gesture to sound as implemented on modern synthesizers.

In the following subsections, I first state a working definition of *musical instrument* which covers both acoustic and electronic instruments. I then address the importance of the mapping from gesture to sound and set that mapping in context as one of four mappings which every musical instrument performs. Following that, I will discuss in detail a second of the four mappings by making a further observation about digital pianos. This second mapping will become the focus of section 2.3.

### 2.2.1 Definition of a musical instrument

A rather broad but nonetheless useful definition of a musical instrument is: a device which transforms mechanical energy (especially that gathered from a human operator) into acoustical energy. A musical instrument thus has two interaction ports; first, a mechanical contact, where transfer of mechanical energy from the human operator takes place, and second, an interface between its resonating body and air, where acoustical energy is transferred into the air. This definition underlines the universality of mechanical input of musical instruments (instruments are generally not voice-controlled or remote-controlled, though electronic technology is providing us with many exceptions). This definition also highlights the fact that the design of musical instruments is in large part a study in impedance matching, both at the mechanical input and at the acoustical output. The piano fits easily into this definition, with the point of contact between finger and key as the human/instrument interaction port and the soundboard as the instrument/air interaction port. The voice does not fit so easily into this definition since the operator and sound production equipment are so keenly integrated. But even in the case of the voice, the drawing of boundaries between instrument and operator can be useful for analytical purposes, though these analyses will not be undertaken here.

Considered from the standpoint of conjugate mechanical variables of force and velocity, and conjugate acoustical variables of pressure and flow, a musical instrument really has two inputs and two outputs. Each port is an input/output pair. A mechanical contact (and an acoustical interface) supports both an input and an output simultaneously. Thus, instruments are much more than devices which transform mechanical input into acoustical output.

A musical instrument, then, is an example of a two-port system as it is known in the field of Network Theory. As network theory stipulates, the signs of both conjugate variables can either be positive or negative, but the designation of one as input requires the other variable at that port to be designated as output. This physical requirement is known as the *principle of causality* in the Bond Graph literature [55]. Thus, at the mechanical port, if force is input, then velocity is output and if velocity is input then force is output. The product of the signs of the port variables indicate the power flow direction. Energy can be exchanged (or power can flow) in either direction, but of course only in one direction at any given time

Note that if we consider both conjugate variables at a port to be capable of carrying information (rather than just their product, instantaneous power), then *information* can flow in both directions simultaneously. One of the mechanical output variables potentially supports information flow back from the instrument to the operator by mechanical (haptic) means at the same time that the operator informs the instrument with the other variable. It is not yet clear from psychophysical studies whether humans can independently monitor force and motion at the same time–though it does seem plausible. This question would have to be answered both for long and short term events and recall times.

There are four possible transformations or mappings from input to output. Each will be discussed below. There exists a mapping (or, for linear systems a transfer function) from each of the inputs to each of the outputs. Figure 2.2 shows our two-port definition of a musical instrument with a mechanical and an acoustical port. Note that the two mappings from the acoustical input are not very important for musical instruments–there is very little energy flow in these directions. I have not labeled the inputs as force or velocity, since that is a decision to be made by the analyst/modeler. Of the four mappings, the one most often overlooked is that from mechanical input to mechanical output (note that both input and output are relayed at the single point of mechanical contact.) This mechanical input/output behavior will be our primary concern in this work, but the mapping from mechanical input to *acoustical* output is also a very important part of what makes a piano a piano and a digital piano not a piano.

Synthesizers and digital pianos, by design, are also musical instruments and fall under the above two-port system definition. We must, however, be careful about terminology when analyzing synthesizers and digital pianos since these systems are not only mechano-acoustical but mechano-electro-acoustical. Analogous to the distinction usually made between the physics of the resonating body and the mechanics of the excitation of that body in the study of musical instrument acoustics,

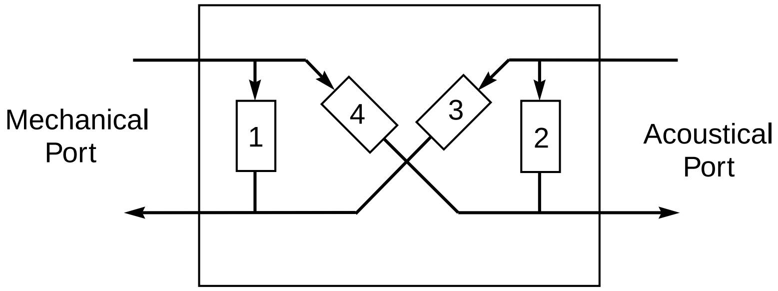

1- mapping from mechanical input to mechanical output
2- mapping from acoustical input to acoustical output
3- mapping from acoustical input to mechanical output
4- mapping from mechanical input to acousitcal output

Figure 2.2: *Musical Instrument As Two-Port System*

there is a distinction to be made between ‘synthesizer’ and ‘controller’ in the design and use of electronic musical instruments. A ‘synthesizer’ is a sound synthesis engine; it takes on the function of the resonating body. A ‘controller’ is a musical interface device (examples of which include a keyboard outfitted with switches, a guitar body with sensored strings, or a clarinet-like tube with pressure transducers) whose purpose is to pick up gestures from the musician. The controller also reduces the gestures of the musician to simple musical parameters which can be relayed to a sound synthesis engine (synthesizer). These parameters are usually encoded as MIDI [4] messages for ‘note on’, ‘note off’, ‘pitch bend’, and so on. Finally, note that the word ‘synthesizer’ in common usage often denotes both synthesis engine *and* controller, since the synthesis engine is often packaged into a keyboard controller, as with digital pianos.

### 2.2.2 Mapping from gesture to sound parameters as performed by synthesizers

Modern digital keyboard controllers typically implement a very simple encoding of the mechanical input (gesture) and likewise a simple mapping of that encoded gesture to the MIDI parameters relayed to a synthesizer. The encoding of input gesture commonly implemented is called ‘velocity sensitivity’. Two switches are placed at successive positions in the travel of a key so that the time between trip of the two switches gives an estimate of the average velocity with which the key was

---

[4] MIDI: Musical Instrument Digital Interface, one of the earliest and most successful communications protocols; it implements standardized communication between synthesizer products of various manufacturers.

depressed. This average velocity is obviously a very crude (in fact, zeroth order) approximation of an input gesture. Neither the mechanical energy relayed by the gesture nor the force input are encoded by velocity sensitivity, and, most importantly, an encoding of the entire gesture (a trajectory over a time interval) is supplanted by a representation with a single scalar. That scalar encoding of the gesture is then typically used as an index in a lookup table to produce the parameter sent to the synthesizer. The user is given access to this lookup table in some electronic instruments. In such cases, the user can set whether the graph of the lookup table is linear, concave, or convex. The output of this lookup table, the MIDI message, then parameterizes the ensuing tone produced by the synthesizer, much like the hammer strike velocity parameterizes the tone produced by a piano string, as discussed in section 2.1. Note, however, that the encoding of the input gesture into a single scalar, right at the outset, cannot be used to emulate the mapping from gesture input to sound output which the piano carries out. Obviously, such a simple input encoding could not implement the same many-to-one mapping. We see that, unfortunately, the paradox of the piano as viewed by an engineer considering the mapping from hammer motion to sound parameters rather than from key motion to sound parameters has been the inspiration for digital pianos. The reduction to a single parameter has been implemented not by a dynamical system such as the piano action, but through the simplistic paradigm of 'velocity sensitivity', a crude representation of the input gesture. The system dynamics of the piano and the interplays between the three other input/output maps of a musical instrument viewed as a two-port system have been disregarded.

This is a grave omission, since the mapping from gesture input to acoustic output is carefully attended to by the musician-users of digital pianos. In fact, with exploratory experiments, it is quite easy for musicians and non-musicians alike to differentiate the mapping implemented on a digital piano from that of an acoustic piano. For example, one can play gradually slower on a synthesizer keyboard and piano keyboard, comparatively monitoring the loudness of the output tone. At slower and slower key depression rates, a synthesizer sounds softer and softer tones, whereas at one point a piano sounds no more tone (the hammer has insufficient kinetic energy to reach the string; no strike occurs.)

In conclusion to this discussion on the mapping from gesture to sound, I would like to note that the correspondence between electronic and acoustic instruments in terms of the excitation means and the sound production means can become quite blurred when considering certain electronic instrument implementations. I suggest that careful scrutiny on the part of electronic instrument designers of each of the input/output pairs of their instruments and the acoustic instruments after which they are modeled will lead to better, more musically expressive designs. Acoustic instruments have one very real advantage over synthetic instruments in that they are implemented by physical means, and thus are subject to some very real restrictions with regard to the available mappings from gesture to sound. Implementation in real physical hardware instead of electronics and software generally promotes the embodiment of dynamic behavior and interesting interplays between the input/output pairs in an instrument. Musical instruments are, after all, designed to be played by humans who are equipped with very high quality haptic and audio sensors, honed for manipulating and gathering information from real physical systems. Because the excitation is effected through physical means on acoustic instruments (a resonating body, whether string, membrane, plate, structure, or air column, is driven in some manner by another body) it can be argued that humans have greater access to the maps of that instrument through intuition, since intuition is generally built on behaviors of the physical world. By association to the behavior of other mechanical systems with which humans have experience, they may already be familiar with the process of excitation as utilized by a particular acoustic instrument. On the other hand, because electronic instruments draw energy from sources other than the mechanical (control) input (they are usually plugged into the wall,) they can implement more arbitrary mappings from mechanical input to acoustical output. These mappings can be very complex or very simple; they are not constrained to those implementable through mechanical design. The energy contained in the control signal does not have to drive the resonating body. The design of electronic instruments usually involves dealing with more distanced abstractions of gesture encoding and sound production than is the case in the design of acoustic instruments.

To summarize, the words ‘impedance matching at the mechanical contact between human and instrument’ do not sufficiently cover the consideration needs of the instrument designer since they imply efficient energy exchange as the only goal and do not take into account the cross coupling of internal connections linking input and output. Designers should also talk of ‘information exchange capacity’ in both directions at the mechanical contact and acoustic interface. Obviously, the information demands of humans are high in both the audio and haptic domains. The degree to which relationships between information appearing on various channels are reflective of physically plausible or organic relationships is also of interest.

In the next section, we will address the mapping in the haptic domain: from mechanical input to mechanical output, but before so doing, we make one more comment about digital keyboards which will lead into the next section.

### 2.2.3 Judged by feel

Digital pianos are judged not only by their sound, but also by how they feel when played. Musicians become quite preoccupied with the feel of an instrument in their hands, and will readily report on differences they detect and shortcomings which they identify. Indeed, synthesizer and digital piano manufacturers are discovering from their outspoken customers just how important the feel of a keyboard really is. K. Yamamoto of Roland, Inc. relates that the design improvement most often requested by customers in recent times is in the area of touch response [108]. High-end synthesizer manufacturers, such as Fanuc of Italy, are even incorporating actual piano actions into their keyboards in order to make them feel right.

Despite the application of assiduous mechanical design, a classically trained pianist will quickly notice that the touch response of a digital piano differs markedly from that of a typical acoustic piano. Thus far, the designs of digital pianos have striven to approximate the feel of the grand piano primarily with the incorporation of passive mechanical components in their keyboard actions. The challenge of design with passive components lies in configuring a spring/damper/mass system which is manufacturable, yet realizes similar dynamics to the piano action. The disparity (as measured either in perceptual or physical terms) however, between the mechanical impedance of these digital piano actions and the acoustic piano action is still quite large. The return force due to the action of gravity on the hammer and key (or static imbalance, as it is known to piano technicians) is comparable, but other factors, such as the balance of spring, damping and inertial component forces, are not consistently emulated. The feel of escapement and hammer bounce (further discussed in Chapter 3), for example, are usually not implemented in digital keyboards.

## 2.3 Dynamics of the grand piano action

In this section, I will briefly introduce the dynamical modeling of the grand piano action. Modeling of the action is an integral part of our approach to the posed problem of emulating the feel of a grand piano with a synthesizer keyboard. However, a description of our approach to the emulation problem and the role of modeling in that approach will be delayed until section 2.4. I will discuss the reduction of a piano action to a mathematical model (a set of ordinary differential equations) which will be suitable for real-time simulation. This section will describe the various dynamical behaviors of the piano action, identify them as the origins of the ‘feel’ of the piano, and implicitly highlight the absence of these behaviors and their associated feels in synthesizer keyboards. Even though the reasons for undertaking the construction of dynamical models will not appear until the next section, we undertake the modeling discussion here in particular because of its relevance in the discussion of keyboard touch-response begun in the last section.

### 2.3.1 The piano as 2-Port system

I have already defined a musical instrument as a system which interfaces with its environment through two interaction ports: a mechanical port and an acoustical port. Each port has two connections, one for each of two port variables. One connection is an input, the other an output. Internally, four mappings relate each of the outputs to each of the inputs. In the last section, we paid particular attention to the mapping from mechanical input to *acoustical* output as implemented on a synthesizer, highlighting its differences with that mapping as implemented on an acoustic piano. In this section, we will turn our attention to the mapping from mechanical input to *mechanical* output, and shift the focus from synthesizer to acoustic piano.

By illustrating the acoustic piano using the definition of musical instrument as two-port, I will clarify what I have been loosely referring to as the 'feel' or 'touch-response' of the instrument. Remember that the mechanical contact between pianist's finger and piano key is the location for energy transmission between the pianist and piano. Certainly the most central function of the piano is to convert mechanical energy input from the pianist to acoustical energy to be output through the acoustical port (soundboard), but also note that the piano action, in addition to having energy conversion and energy dissipation elements, has numerous energy storage elements. The piano thereby has a capacity to store energy input from the pianist and return it back to the pianist through the contact at the key. Also in the last section, we suggested that both port variables are capable of carrying information, that is, capable of being monitored or tracked by the player. We will therefore center the discussion on the origin of the feel of the piano around the port variables of force and velocity rather than their product, power.

### 2.3.2 Mechanical Impedance

It is the relationship between the two mechanical port variables, force and velocity, which give rise to the 'feel' of the piano at the key. If we identify the velocity as input and force as output, the feel is characterized conveniently as the mechanical driving point impedance [5].

Mechanical impedance defines the dynamical (history dependent) relationship between force and velocity, with velocity in and force out. Admittance also describes a relationship between force and velocity, except that the force is input and velocity output. In laymen's terms, admittance can be described as the 'give' of a mechanical system. For a given velocity input over time, mechanical

---

[5] In its strictest sense, the word impedance pertains only to linear systems, expressing the pertinent mass, damping, and spring component forces in response to velocity input. In this work, however, we use the term 'impedance' to refer to the force-velocity relationship of 'nonlinear' systems as well. Nonlinearities may be due to nonlinear constituent laws of spring and damping elements and system discontinuities such as those caused by making and breaking of contact between bodies internal to the piano action.

impedance can be used to express how that system will respond with force. Impedance is also conveniently expressed in the frequency domain and is thus occasionally defined as the frequency dependent mapping from velocity to force. For example, to a certain input motion (velocity), the piano key will respond with a certain interaction force, and do so in a way which is dependent of the history of the input velocity. Note that the definition of mechanical impedance is completely analogous to electrical impedance, where, instead of force and velocity, the port variables are voltage and current.

#### Vibration Output

Driving point impedance does not fully account for the feel of the piano at the key. The piano is actually a multi-input, multi-output system if we define each key and the soundboard as a port. In addition to the response from the input at that key, there is response from inputs at other keys, and there can be response from energy put into the piano quite some time ago. The piano is, by design, capable of energy storage in the form of vibration of the strings and soundboard at acoustical frequencies; and these vibrations can be felt at the keys. Measurements of typical vibration frequencies and amplitudes at the key of the piano were made by Askenfelt and Jannson and shown to be well above the sensation thresholds for humans [4][6].

Yamaha and other synthesizer manufacturers have attempted to duplicate the vibratory haptic cues of the piano in their digital piano designs by mounting the bass speaker on the keyboard frame in such a way that it produces vibrations at the keys typical of a grand piano.

In this work, however, we will be interested primarily in that portion of the driving point impedance which is definable for a single key; what would be described by the word ‘give’ rather than ‘vibration response’.

### 2.3.3 The modeling challenge presented by the piano action

The piano action is a compound lever system supporting a myriad of performer/system interaction behaviors. Our goal is to capture in dynamical models those behaviors which are salient for the mechanical driving point impedance at the key. Especially since we will later require these models for real-time simulation, we will construct models which are as simple and computationally efficient as possible, yet preserve the behaviors which we have targeted for emulation. This task will not be

---

[6] For a particularly interesting discussion of vibration-range haptic response from stringed instruments and its utility, see [4] and other articles in the Spring 1992 issue of Music Perception, devoted fully to Somatosensory Feedback in Musical Performance. In this reference, the use of vibratory cues from the cello, double bass, a singer’s chest, and the trumpet are experimentally characterized and compared to human sensation thresholds.

simple since many of the behaviors which the piano action supports are due to effects which are difficult to model with present-day engineering tools.

At first glance, the piano action is like many other mechanical systems in that it stores both kinetic and potential energy, dissipates energy, and transmits energy to other mechanical systems. But there the simplicity ends. Among its component parts are those which must be described by nonlinear constituent equations to capture their effects, most notably the felt and leather couplings. Furthermore, it is a dynamical system which undergoes changing kinematic constraints (bodies make and break contact with one another). It is thus an inherently discontinuous system. In fact, the particular kinematical constraints operative at a given time are not only a function of its configuration, but also a function of the history of the configuration up until that time. For example, depending on how the key is hit, the hammer can be caught on the backcheck or on the repetition lever. The piano action features both holonomic and non-holonomic constraints, especially sliding constraints.

Owing to the non-linear and discontinuous nature of the piano action, it is easy to recognize that inputs and outputs will not scale in linear fashions.

### 2.3.4 Models of the grand piano action

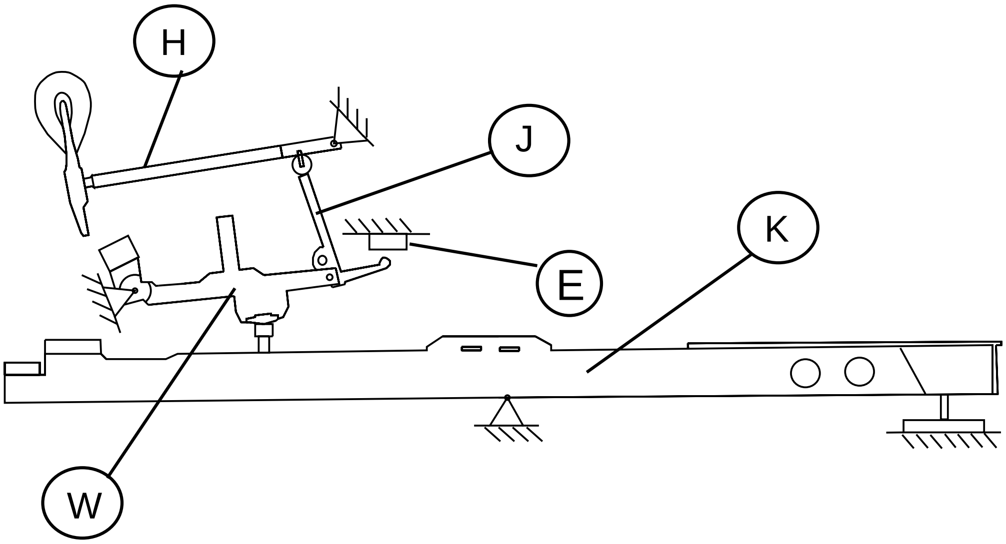

Figure 2.3: *The Elements of the Piano Action*

Figure 2.3 shows a somewhat simplified model of the piano action. It is an assembly of four bodies, known commonly as the key $K$, whippen $W$, jack $J$, and hammer $H$. The escapement dolly $E$ is fixed to ground. There exist many excellent references covering the history, design, maintenance, and regulation of the piano action. See, for example, [84]. Rather than introducing the behavior of the piano action with reference to Figure 2.3, I will introduce its behavior with a simplified model, one which is rather like a ball bouncing on a paddle. This model will be indicative of those described more fully in Chapter 3.

Beyond the static return force due to the action of gravity on the hammer and key, the feel at the key is dominated by the inertia of the hammer, over which the pianist has a five-times mechanical advantage from the key. The piano action can therefore be modeled, to first approximation, as a static return force along with a simple mass sized to match the effective inertia of the hammer as felt at the key. Because the entire swing of the hammer is less than 25 degrees and that of the key is less than 5 degrees, it is reasonable to approximate this system of rotating levers with a linear system. We assume that all interaction forces and gravity forces act perpendicular to the bodies to which they are applied. We also neglect the effect of sliding between levers.

Figure 2.4 shows a simple linear model for the piano action. The mass $m_h$ is a linear approximation to the rotary inertia of the hammer, scaled by the mechanical advantage squared, or 25. The additional mass $m_k$ represents the inertia of the key, which has a magnitude about 50 percent of the scaled inertia of the hammer. The stiffness || of the connecting spring represents the stiffness of the wooden levers and the felt and leather covers at the various contact points between key, whippen, jack, and hammer. The jack and whippen are not accounted for in this model since their inertia does not contribute significantly to the feel at the key.

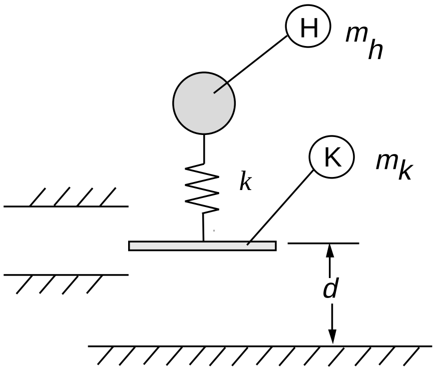

Figure 2.4: *Model 1*

The double-mass model of Figure 2.4, however, is representative of the piano action in only

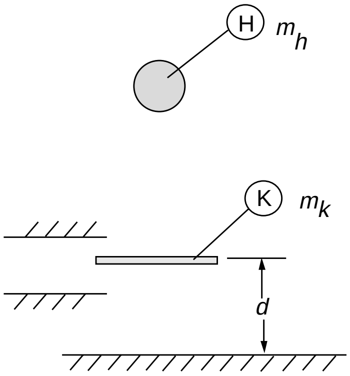

Figure 2.5: *Model 2*

one of its configurations, the initial configuration in which the hammer rests on the jack. This configuration is abandoned at a certain point in the motion of the action —another configuration called 'letoff' takes over. Letoff is the second phase in a sequence of phases or periods of configuration of the action. Letoff begins when the jack meets the escapement dolly $E$ in Figure 2.3 and ends when the hammer flies free of the jack. During letoff, the jack pivots out from under the hammer knuckle where it has been pushing, allowing the third phase to begin, which I shall call 'free-flight'. Here in this introductory section, we will construct a model comprising only two submodels, one for the coupled phase and one for the free-flight phase. Models which include the letoff phase and thereby also capture the effect of letoff in the feel at the key will be presented in Chapter 3.

During free-flight, we represent the piano action by a model in which the hammer is decoupled: a simple mass falling under the action of gravity towards the coupled configuration. Figure 2.5 shows our model for the free-flight phase.

The incorporation of limit stops for the motion $K$, as seen in Figures 2.4 and 2.5 gives to our model the function of the front and back keybeds (keyrails). The body $K$ is constrained from traveling beyond the upper limit stop, causing $H$ to decouple from $K$, given that it has enough kinetic energy for lift-off (the spring force shall never be tensile).

This simple 'bouncing ball' model of the piano action supports various interactions between pianist and piano. If the key is pushed down gently, the hammer may never leave the key. If a stronger strike is made, the hammer will leave, and a change in inertia of the key can be detected both when the hammer decouples from the key and later when it lands back onto the key. With a model made up of two submodels we aim to duplicate the effects of the changing kinematic constraint in the piano action arising from the hammer 'bouncing' on the jack. By sequencing back and forth from the coupled to the free-flight submodels as a function of the input, we realize the dynamics of a discontinuous system. For a more detailed description of this two-body model of the piano action, see [33].

Note, however, that the simple model comprising the submodels of Figures 2.4 and 2.5 cannot even be called a piano yet, since it does not have an escapement mechanism. The hammer must be prevented from re-striking the string under the action of a single keypress with an escapement or letoff of some kind. An escapement was the central feature of Bartholomeo Christofori's original invention of the piano in 1705. Other functions of the piano action which will be defined and modeled in Chapter 3 are: letoff and effects of the repetition lever.

### 2.3.5 Form of the Model

Mechanical system models come in many forms. Because we are interested in running real-time simulations of models such as the one introduced above, we will favor models expressed in forms which lend themselves to efficient simulation. A model can be expressed as an ordinary differential equation (in which the constraint equations are incorporated), a differential algebraic equation (the algebraic constraint equations are adjoined to the differential equations), or a set of coupled second order systems. Each model expression will offer certain advantages for real-time simulation.

Of the various forms in which mechanical system models are expressed, a set of ordinary differential equations (ODEs) in which the constraints are incorporated (used to eliminate dependent coordinates), rather than adjoined, is the simplest. A set of ODEs with incorporated constraints shall be a 'model in independent coordinates', or alternatively, a 'reduced model'. The full model is then formulated as a piece-wise continuous ODE. Discontinuities are allowed at timepoints corresponding to changes in the kinematic constraints. The time periods between the discontinuities are each governed by one of a set of 'submodels', each of these being a continuous ODE constructed to describe the system in one of its constraint conditions.

In the next section, we will address and motivate our approach to real-time simulation of these mechanical models, paying particular attention to the realization of discontinuous systems, *i. e.*, those comprising submodels.

## 2.4 The Touchback Keyboard

In this section I introduce our solution to the problem of emulating the feel of a grand piano in a synthesizer keyboard which does not include whippens, jacks, and hammers. We do it with motors. Each key is driven by a dedicated motor with which interaction forces are fabricated, under computer control, as a function of the monitored motion of the key. In this manner, a programmable mechanical impedance is created at each key. I will discuss the architecture and algorithms in this section which lead to the emulation of a piano-like mechanical impedance with motors, sensors and computer.

Computer mediated emulation of mechanical impedance is known today as ‘haptic display’. This field is quite new; dating back in explicit mention only to about 1980. For a very complete review of haptic interface technology back to its origins, see [71]. Here, I will relay only a few anecdotal notes to place haptic interface technology in perspective.

Robots are generally designed to execute tasks or operate on their environments much like a human and generally in place of a human [26]. To have a human touch or interact with a robot arm is expressly avoided; such activities are (appropriately) regarded as very dangerous. To in fact design a robot, however, for the purpose of interfacing to a human through a mechanical contact is exactly the goal of haptic interface. Haptic interface devices are actually robotic devices– motorized manipulators. But rather than being purposed for manipulation, a haptic interface is purposed to be manipulated by a human and is therefore often called a ‘manipulandum’.

Mark Bolas, president of Fakespace Labs, Inc., a firm developing boom-mounted visual displays for application in virtual reality, places the phenomenon of haptic display and virtual reality in an interesting light [12]. Since robots have not been able to enter into our world and take over our chores as was expected back in the ’50s, we are now taking it upon ourselves to enter into their world. We are the ones clothing ourselves in their wares: we wear head-mount displays, donn sensored gloves and put on haptic displays in order to interface to the worlds of the computer.

In this section, I will briefly review existing applications of haptic interface technology to music. I will provide a framework description of the basic activities involved in the creation of a touchable virtual object which is modeled after a real-world object. Our own methods will be placed in this framework. Dynamical modeling plays an important role in our methods, and therefore this section provides the motivation for the dynamical modeling of the grand piano discussed in the last section. In the present section, I highlight our formalism and system architecture, especially the manner in which we formulate our models for real-time simulation and incorporate them into the haptic interface hardware and software.

### 2.4.1 Haptic Interfaces in Music

At Stanford’s Center for Computer Research in Music and Acoustics (CCRMA), it was Max Mathew’s idea in late 1988 to address the issue of the lacking touch response of synthesizer keyboards with haptic interface technology. Although Professor Mathews was at the time not aware of the then fledgling field of haptic interface, his original sketch, as seen in Figure 2.6, can be recognized as a call for haptic display in music. Professor Mathew's sketch includes a motion sensor, a digital controller, and an actuator in the form of a solenoid.

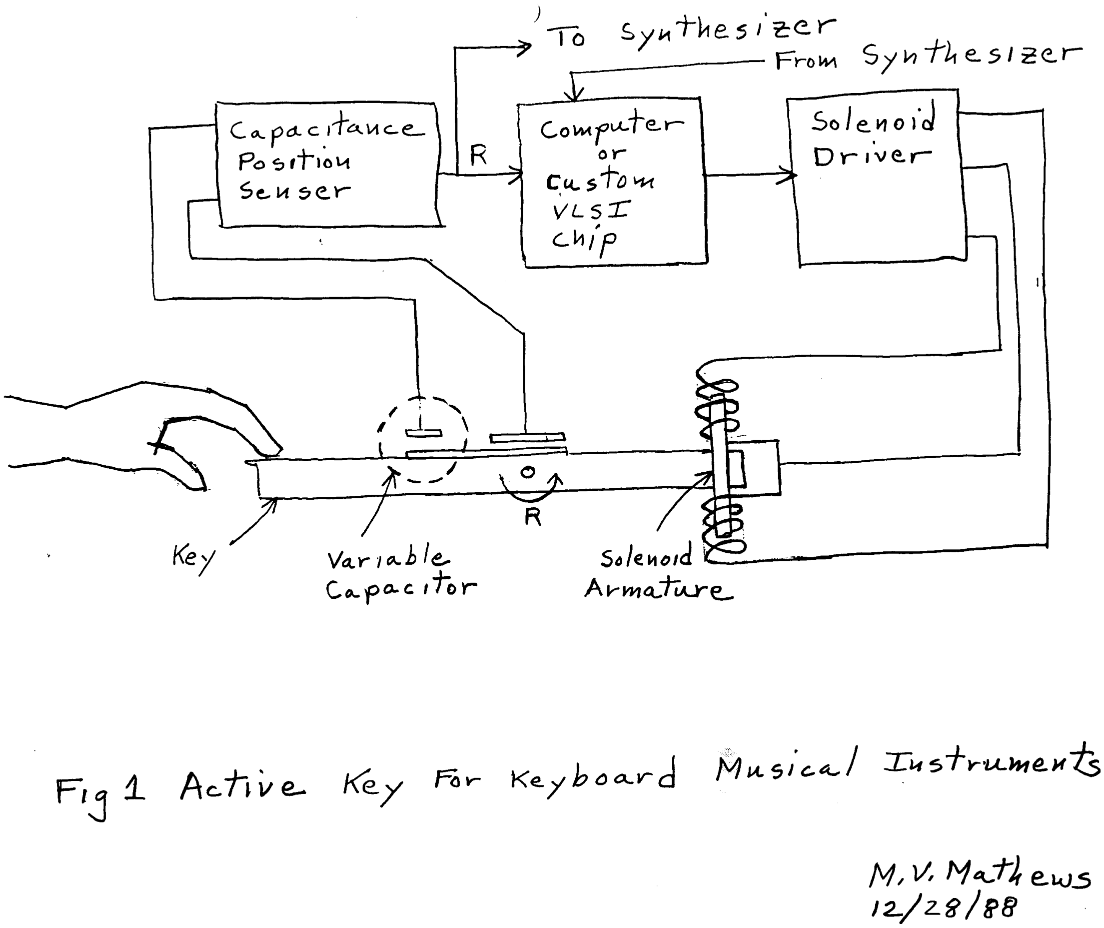

Figure 2.6: *Sketch of Active Key idea by Max Mathews*

Claude Cadoz' group, Association pour la Création et la Rechereche sur les Outiles d'Expression (ACROE) in Grenoble, France began developing haptic interfaces for emulation of virtual musical instruments in 1984 [19]. By 1990, Cadoz' group had constructed a motorized 16-key keyboard of a patented stacked-magnet design and an extensive software support library. The controller software was developed to run on transputers [18]. I will make more comments in Chapter 3 about their simulation algorithms, highlighting our methods and theirs in a common framework. Certainly our work owes much to this early work at ACROE in the area of haptic display for musical application.

Richard Baker holds a US patent for a force-reflecting keyboard [5] which he calls the 'Active Touch Keyboard'. His design includes keys coupled to small motors through a cable and pulley arrangement. An analog controller facilitates the setting of static imbalance and resistance forces. Unfortunately, Baker's design has not been commercialized to date.

Baker's keyboard is not able to emulate the touch-response features of letoff resistance (the rise in reaction force during the letoff phase due to extra friction) or the bounce of the hammer on the jack. The effects of changing kinematic constraints are not captured by the analog controller. We hold features which are localized in the stroke of a key and associated with changing constraints to be central qualities in the touch-response of the piano and primary targets for emulation. Baker's approach using an analog controller, however, is intriguing. An analog controller is not prone to the destabilizing effects of discrete control. By the way, compensation for the destabilizing effects of discrete control occupies a significant portion of this thesis (Chapters 5 and 6). We have used digital control because of its comparatively superior (to date) programmability. Emulation of discontinuous effects (due to changing kinematic constraints) with an analog controller would be possible with nonlinear switching elements such as diodes. An implementation with an analog controller is actually an extension to the work of this thesis.

Alistair Riddell designed and built a 'Meta Action', a set of solenoid-driven hammers to hit strings in an otherwise complete piano. The meta action was designed to take the place of the action and realize manufacturability advantages [85]. We regard the meta action to be an incomplete solution, since it preserves neither the touch-response nor the intricate mapping from gesture to sound parameters of the piano action.

We have designed and built a keyboard-like haptic interface with eight keys which, with its controller, is able to emulate the feel of various keyboard instruments. Interaction with a virtual piano action as is made possible with such a motorized keyboard is a very promising means of re-establishing the kind of relationship between musician and keyboard which supports sententious musical expression. In fact, we believe that haptic display technology will one day become a more viable means of creating the desired touch-response in commercial instruments than mechanical design with passive components. The most notable advantage is the intrinsic programmability of the virtual piano action. The feel of a harpsichord, piano, forte-piano or some altogether new instrument would each be available at the push of a button. But further, the virtual action can be programmed to suit personal preferences. Thus the relationship between touch-response and expressive control can be explored on an individual basis.

### 2.4.2 System identification, perceptual modeling, and physical modeling

Having posed the basic problem as one of emulating the mechanical impedance of a real-world device with a simple motorized device, we must start with a representation of the impedance of that device, one which will be useful as the core of a controller for a haptic interface. The challenge lies in representing the target device's interaction physics in some form or expression which can be utilized as the controller for the haptic interface. The controller will cause the interface to behave in the user's hand in such a way that the user is led to believe that he is interacting with the target device (when in fact he is just interacting with a motorized manipulandum). By choosing a particular representation, we are in large part also choosing an architecture for the haptic display controller. Likewise, by choosing a method for characterizing the impedance of a device for emulation, we will be favoring certain representations. Since the characterization or modeling process will have such a large influence on the design of a device's implementation as a virtual touchable object, I have identified three distinct methods for characterization.

I propose that the creation of touchable virtual environments encompasses two basic activities: first, a characterization or encapsulation of the target impedance, and second, an incorporation of the characterization data into the software and hardware design of a haptic interface. The first of these activities, characterization, can proceed according to one of three basic paradigms:

1. system identification,
2. perceptual modeling,
3. physical modeling.

Each of these approaches to characterization will later favor a certain design architecture in the haptic display hardware and software.

First a brief note on each method to highlight their differences: The first method involves *system identification* of the target device, and subsequent implementation of the reduced model into the controller of the haptic interface. The second method, *perceptual modeling*, involves assuming a model and then iteratively tuning model parameters for optimal emulation as directed by human subjects making comparisons through haptic exploration. The third method, *physical modeling*, involves building a model of the target system, incorporating that model into a simulator, and running real-time simulations with the human in the loop via the haptic interface.

Let us now consider these characterization methods one at a time.

#### System Identification

*System identification* is described by the following procedure. The target system's impedance is characterized empirically, that is, data for the mechanical impedance are collected from the physical device itself into a suitable data structure. That data structure might be a simple lookup table or even a multiply indexed lookup table. The collected data must then be implemented as a formula, algorithm, or lookup table (as befits the chosen data structure) in the controller of the haptic interface. Note that the system identification process can be conveniently accomplished using the haptic interface itself as the probe, if it is instrumented with both force and motion sensors. One can even imagine the ultimate device which would itself feel and characterize the target system, then turn around and present what it felt.

### Perceptual Modeling

The second characterization method, *perceptual modeling*, depends on a human subject to make judgements about the quality of whatever attempts are made at emulation. It is with reference to an already rendered virtual object, and its comparative evaluation against a real object or a perceptual (mental) model of such, that a suitable representation for the real object is found, that is, the real object is modeled. This is typically a trial-and-error approach, though measures can sometimes be established for more efficient searches [70]. The use of perceptual models can side-step many difficulties of system identification and physical modeling which depend to a greater degree upon engineering judgment for their success. One may happen upon (or even be directed to by engineering intuition) an algorithm which is found interactively by touch to render a desired effect, such as a rough surface or a particularly hard wall. The selection of the model effectively takes place through evaluation of the rendered effect. An altogether different activity also constitutes perceptual modeling according to this definition: the selection of certain real objects for the purpose of practicing a manipulation task. Surgeons, for example, practice epidural analgesia with a needle through tomatoes and pears, these being their perceptual model for the biological tissues encountered by the needle during that procedure [34].

### Physical Modeling

The third characterization method, *physical modeling*, is described by the following procedure. The target dynamical system is modeled in the engineering sense, that is, reduced to a mathematical description such as a set of differential equations. Parameters in the model are estimated from known physical parameters of the system or derived from separate experiments on that system. For example, the moments of inertia and mass properties may either be known or easily determined from the target system with observations of dynamical behavior in response to known inputs. The model is then typically used to formulate an impedance expression (force output is expressed as a function of input motion). Note that a physical model can have many forms, and each form may favor a certain implementation in the haptic interface. Comments on the implementation of a physical model into a haptic interface will be made shortly.

### Our Approach

Our approach to the rendering of the impedance of the piano action with haptic display has been inspired by physical modeling, the third characterization method introduced above. We have concentrated on the physical modeling paradigm for the following reasons:

Firstly, we consider the relatively straightforward extension of a model by simple variation of model parameters a significant feature. Also, the manner in which a model grows in complexity is due to an engineer's application of modeling judgement. This implies that accountability is always present in the process, and explosions in complexity are more easily avoided. By the same token, model development is not so easily automated, at least not to the extent possible with either the system identification or perceptual modeling methods. This fact, however, we consider to be a feature rather than a detractor, since at this early stage in the development of haptic interface technology, the incentive for automation is low. Other advantages of physical modeling include the fact that models can be succinctly expressed, shared, and published. Various modeling methods can be applied side by side or by independently working researchers to the same system to confirm results.

The suitability of the particular approach taken depends to a large extent on the nature of the impedance to be rendered. If the impedance is easily expressible as a lookup table, then a system identification approach is straight-forward. For example, multiple lookup tables, one indexed by position, one by velocity, and one by acceleration could be used with their outputs summed to implement a non-linear but superposing impedance. Discontinuous systems (those whose impedance takes drastic jumps as a function of configuration or configuration history), on the other hand, will likely be better rendered with the use of a physical model. Because we felt that the piano action fits into this latter category, the physical modeling approach distinguished itself from the outset of this project.

Finally, very powerful, flexible, and extensible dynamical modeling and analysis tools were available to us early in the undertaking of this project. In particular, the dynamical system analysis program *AUTOLEV* [89] for the generation of equations of motion by automated symbol manipulation presented itself as an extremely useful model construction tool.

### 2.4.3 Implementation of a physical model

Whereas the implementation of the data structure produced by a system identification activity is relatively straightforward, and a perceptual model depends in any case on a previously cultivated implementation for its production, the implementation of a physical model into a haptic interface is not so simple; it deserves more comment.

Implementation of a physical model takes the form of a real-time simulation of that model. Certain control variables are fed into the simulation in real time, directly from motion sensors on the haptic interface. Examples for the physical modeling approach include flight simulators, although rather than haptic display, flight simulators use motion display.

For each basic form of a physical model, as discussed in section 2.3.5, a corresponding haptic interface implementation can be suggested. If the model is static (and algebraic expression) it may be implemented directly as a control law. If the model is in the form of a differential equation, and if a solution as a function of the input (sensored) variable can be found, it can implemented directly as a control law. If it is a linear differential equation, it can be converted to a discrete formula, and the solution at each time step is obtained by a simple matrix multiplication. If the model is an ordinary differential equation, it can be wrapped with a numerical differential equation solver such as a fourth order Runge-Kutta solver. These implementations will be more thoroughly discussed in Chapter 4. It is also possible to implement either an admittance or and impedance controller for haptic display. These two basic approaches will also be compared and contrasted in Chapter 4.

### 2.4.4 Our System Capabilities

I will now turn from the discussion of haptic interface implementation in general to a discussion of our implementation in particular.

We obtain the equations of motion governing the behavior of the grand piano action using $AUTOLEV$, paying particular attention to the accommodation of changing kinematic constraints. We have chosen to formulate the equations of motion in their reduced form (incorporating the kinematic constraints) so that they may be integrated by a standard ODE solver. Thus, the rendering of a system with multiple constraint conditions requires the formulation of multiple submodels and a passing of the state information from one submodel to the next at the transition times. Each sub-model governs the behavior only during that time-period which corresponds to the particular kinematic structure for which it was developed. Thus, our simulator is able to handle models described by ODEs which are piece-wise continuous, with the discontinuities occurring at times which are themselves functions of the state. In this manner, a simulation of the piano action is able to account for the changing kinematic constraints that occur when the elements of the action make and break contact with one another.

The rather simplistic model of a bouncing ball developed in the 2.3.4 comprising the two submodels of Figure 2.4 and 2.5 has created a very convincing virtual bouncing object when implemented with a motorized key. Interaction between the ball and user through the key (in this case to be viewed as a paddle handle) includes all the properly timed power exchanges to suggest manipulation of a real ball and paddle. As suggested in the Section 2.3.4, discontinuous (but otherwise linear) bouncing ball model can be thought of as the hammer bouncing on the jack. In summary, we have implemented a unilateral constraint (a gross non-linearity: a contact capable of supporting compressive but not tensile forces) by combining two linear submodels with some management routines for exchanging them in and out of the simulator.

In the most general rigid-body mechanical system, the various constraint conditions may be taken on in any order, depending on how other systems (possibly a user) interact with it. Barzel [9] has addressed the realization of discontinuous systems by simulation of a sequence of ordinary differential equations. In our work, we adopt his nomenclature and combine it with a Finite State Machine (FSM) simulator, which will allow a sequence of conditions or 'states' to be taken in an order which is not known ahead of simulation-time.

## 2.5 Optimization

We have been working under the premise that the touch-response or feel is a crucial component of the piano, intricately tied to its capacity as a musically expressive instrument. We have been developing means to incorporate the feel of the piano and other keyboard instruments into synthesizer keyboards.

The role of the 'feel' of a keyboard instrument in the music making process is rather subtle, as has been underlined by the discussions in sections 2.1 and 2.2 above. Verily, the mechanical energy exchanges between fingers and keys which give rise to the 'feel' cannot be treated separately from the process of converting from intentions to sound output without losing sight of the investigative purpose. The pianist, who has certain objectives in mind when manipulating the instrument, is subject to the behavioral features or the 'dynamics' of the piano action. The pianist must operate within its constraints, and utilize, to the best of his or her ability, what information about its response that it makes available. We are ultimately interested in the extent to which the pianist has control over the piano. Our measure of control is the degree to which certain objectives are met, especially certain musically significant objectives. Also of interest is the degree to which the pianist can vary the piano's output along chosen lines, that is, lines deemed musically significant.

In particular, we ask: how will the pianist's attempts at varying a single parameter independent of others be met with success, and how are such relationships between mechanical input and audio output dependent upon the relationships between mechanical input and mechanical output of the instrument?

Also note that the control inputs which a pianist uses are the product of a very long development period. Their form is by essence the product of much practice. In fact, there are few human endeavors which enjoy as much devotion and time commitment as piano practice. A professional pianist typically spends three practice hours per day at the keyboard; the professional pianist in training will spend as many as eight hours per day.

### 2.5.1 Human as Optimizing Controller

Many investigations into human performance have been based on the premise that the human is an optimizing controller. For example, see studies of pilot behavior by McRuer [66]. This leads us naturally to address piano performance with optimal control theory. It seems quite plausible that the purpose of practice is to optimize the piano's response, to bring it as close as possible (according to some measure) to a chosen objective. The pianist strives for optimality despite given constraints and in the face of given detracting influences. In other words, the pianist must work within the bounds posed by performance variations, repeatability, variations between keys on the keyboard or between pianos, and attempt to express a musical interpretation for a particular audience and time.

Indeed, it seems plausible that the mechanical impedance of an object provides the human manipulating that object with a great deal of information as to its behavior and its variety of behavior. It allows the human to develop an internal model or 'intuition' about the system being manipulated. In most every sport, it would be unthinkable to deny the athlete haptic interaction with the sporting equipment. That would be tantamount to robbing the athlete of the source of satisfaction of the game.

Optimal control is a well developed field. There exists an extensive toolbox to handle all sorts of constraints: nonlinear constraints, even model changes (changing kinematic constraints). Optimal control theory has a long history and many proponents. It owes its roots to the calculus of variations, and found extensive application in flight mechanics during the development of aeronautics technology. Today it enjoys a wide application area including economics, mechanical design and process design.

### 2.5.2 Applications of Optimization Theory: Motivation for Haptic Interface

We have undertaken an optimal control analysis of the pianist/piano system dynamics primarily to elucidate the role of the feel of the piano. The aim of our optimal control study is to give a solid basis for the claim that the feel of a keyboard instrument is important to preserve in electronic keyboard designs. An approach which considers the pianist as optimizer, subject to certain constraints and system dynamics, lays out what we think are the proper roles for each of the players in the game: piano/pianist/musical objective. A block diagram analysis of the pianist as controller, the piano as plant, with both haptic and audio sensors feeding back signals to the controller, is a good starting point. But, if the questions we are interested in asking have to do with the value of feedback, it is a sensitivity analysis or an optimality analysis which we must undertake.

Our intention in using optimization theory is not to develop an optimal control input or find the performance limits, but rather to make certain points about the performer/instrument relationship. We wish to make the role of the feel of the piano explicit in order to motivate the application of haptic interface to this problem. We are not particularly interested in developing a player piano or, as is often the purpose in optimal control studies, an autopilot. We do not plan to use an optimized path in the control law.

We are stepping back to try and answer the questions raised above by applying tools from dynamic optimization theory. We make the plausible assumption that the human control input, after all that practice, is in fact optimized either for efficiency or its effectiveness at attaining some musically significant objective.

By imposing quadratic objectives, and using simple models of the human finger playing a piano or synthesizer action, the optimized control inputs are found. Not surprisingly, the optimized inputs are very much a function of the dynamics of the action–which suggests that piano technique will not transfer to a synthesizer keyboard, as is observed. Various more musically significant objective functions are being studied at present. Treatments of motor noise or lack of repeatability, for example, are current goals. Future experiments are foreseen in which human subjects are asked to repeatedly play their best pianissimo at a force and velocity sensored keyboard. These inputs could then be compared to the analytic ‘robust pianissimo’.

## 2.6 Summary

Humans are admirably equipped to explore and characterize the mechanical properties or behavior of objects in their environment. Our muscles and articulated limbs allow us to manipulate, and our haptic senses provide us with information about an object’s mechanical response to our manipulations. Our goals, however, often go beyond system identification or characterization of the mechanical impedance of an object. We may want to influence an object’s behavior. Such is the case when playing a musical instrument. While attempting to exact a desired dynamical behavior (and corresponding sound) from an instrument in our hands, we use not only the sound but also the force/motion response in a feedback sense to modify our manipulation. If the musical event is of a long duration in comparison to human response times (typically greater than 200ms), this information may be used by the musician for real-time feedback control. Otherwise, the response information is used for anticipatory control to modify the manipulation the next time around, as with practice and learning. Examples of musical instrument playing in which haptic information is of relatively obvious and immediate value to the player include 'sultando' or 'ricochet' bowing of a string instrument and use of the repetition feature on a grand piano. When an instrument's mechanical behavior has no correspondence to its acoustic behavior or there is no haptic information available, a very valuable channel of communication from instrument to player is lost. This is the case for most synthesizer-based musical instruments. In order to alleviate this deficiency on keyboard synthesizers, yet preserve and even expand their programmability, we are developing a synthesizer keyboard with haptic display.

In the foregoing chapter, we have studied certain psychophysical phenomena underlying the operation and design of the piano. The paradox of the piano, that players seem to have independent control over timbre and intensity, yet the percussive nature of the instrument does not support such independent control, was legitimated rather than denied by acknowledging the involvement of psychophysical and psychoacoustical factors. The fact that players only have control over two scalar parameters for each note, hammer strike time and hammer strike speed (the sound parameters) was discussed. In support of fine control over these parameters, the importance of features in the mapping from mechanical input to mechanical output at the key (impedance) and the mapping from mechanical input to acoustical output (gesture to sound parameters) was highlighted.

We are interested in exploiting these psychophysical phenomena in the design of future electronic instruments, so we have considered their application to date in synthesizer controllers on the market. The unsatisfactory level of control over sound parameters which synthesizer keyboards make available may be attributed to the fact that their design does not include the features of the mappings which are characteristic of the piano.

The mechanical impedance of the piano is mediated by its physics or 'dynamics', so we have studied these by building simple dynamical models of the grand piano action. These dynamical models can be simulated in real-time in a human-in-the-loop control scheme with display through a haptic interface. Our Touchback Keyboard, a prototype haptic interface specialized for the emulation of the feel of the grand piano action, was briefly introduced. Further studies of the principles of control which a human uses can be undertaken with a keyboard-like haptic interface such as the Touchback Keyboard. The observations made in this chapter about the piano may of course be extended to other instruments.

All musical instruments are manipulated through some kind of motor (muscle) control, usually involving a mechanical contact. Through a physical contact, a musician guides the mechanical behavior (and in turn, the acoustical behavior) of the instrument. Even the most contrived and out-of-real-time of electro-acoustic instruments has a some kind of mechanical interface. It is the effectiveness and appropriateness of this mechanical interface with regard to a musician's desire to express himself which concern us here. The haptic senses, by picking up information about the mechanical behavior of the instrument, can glean information about the acoustic behavior if there exists a relationship between the acoustic and mechanical behavior. In most acoustic instruments there does indeed exist a close correspondence between the touch response and the sound response.

The goals of a musical instrument designer may be succinctly stated using two concepts from modern controls theory: controllability and observability. The goal is to maximize controllability and maximize observability. Maximized controllability in an instrument suggests that the performer has very fine control over the instrument's output (and perhaps numerous ways to exercise that fine control). Maximized observability suggests that the instrument makes available to the performer a maximum amount of information about its behavior. And of course one mode by which the instrument may inform the performer is through haptic stimuli. Maximized controllability and maximized observability go hand in hand to enhance the relationship between performer and instrument.

It is not surprising that humans happen to be well equipped for listening to and playing musical instruments since most instruments were invented by humans to be thusly listened to and controlled. As we invent new instruments, however, especially computer-based instruments, we must carefully analyze the musician/instrument relationship. With the computer, after all, we are re-defining the way in which music is composed and performed. In particular, new computer-based music interfaces must consider the existence of mechanical information exchanges in both directions between performer and instrument.