Loose Ends — Index

0. Introduction

2001: Early this year, I went through my collection of books, articles, correspondence, notes and other artifacts with the goal of reminding myself what projects I was in the middle of. The notes I took ended up being roughly divisible into three categories: *principles* worth remembering or pondering, *projects* contemplated, started, or completed, and *resources* that might come in handy at some point. It struck me that, but for the fact that the items on the list were in the form of terse notes to myself, this list would be a fairly complete answer to the question "what are you working on?" — a question I often come up short on — so I undertook to add enough detail to the items on the list to make it more comprehensible to the general reader. This document is the result.

2002feb23: From last year until now, I've been adding and updating things.

2005jan09: Starting with this version, I'm going to create dated versions when I edit this document.

1. Principles

I think of these "principles" as akin to mathematical "operators" in the sense that they can be applied in interesting ways in a lot of situations (though some of them are specific to teaching, making presentations, music, etc.). My intent was to use these as general-purpose stimulants — ideas to provide the inspiration to get over (or a distraction to sidestep) a conceptual hump. Or perhaps "conversation pieces" is a better description. They are not rules.

1.1 General

1.1.1 General principles and observations

Context. Perhaps the most universal conceptual operator. There are usually many contexts in which it's appropriate to consider an idea. Once an idea has been explored within a certain context, it's often useful to abstract it from that context, and then see whether it can be applied in other ways in other contexts.

Proximity. Visual proximity suggests conceptual proximity. This seems obvious, but failures to use the principle abound. (For example, what "order" is there in "alphabetical order"? None. Imagine how much easier it would be to learn and remember the alphabet if it were in some kind of rational order, for example, if all the vowels were first.)

To kill two birds with one stone, you have to see two birds close to each other. (This is related both to "context" and "proximity.")

Use a dynamic presentation to show a dynamic entity (such as sound/music). (This is not always appropriate, however; a musical score, for example, is static.)

Individual results are sometimes less important than the perspective/approach/method/tool. For example, it is valuable to know what it is to be an expert in something (even something useless), because it gives you a point of reference : when you're learning something new, you can realize "I know some things about this, but I'm not an expert." Likewise, "knowing how to do something" (or how not to do it) is more valuable than any particular result.

Ubiquity. One of my favorite examples of ubiquity is language. It's always with you. So, when designing something, it's worth considering: how can we make this more ubiquitous?

Mathematics. What is mathematics other than a substitute for a perfect imagination (and memory)? Mathematics is a language for communicating in condensed form ideas which would otherwise be too unwieldy; it is an economical and convenient way to remind ourselves where our thoughts led. Why is the language, the notation of mathematics the way it is? Is it because of our history of (one dimensional) spoken, written, and printed language? The notation of mathematics has gotten somewhat less linear (arrays, etc.), but the question remains open: could there be a better way? Conventional mathematical equations are, inherently, one-dimensional, across the equal sign. Nature is multi-dimensional. Could we develop a new form of mathematical expression that mirrored this?

1.1.2 Approaches

Show the space. That is, when there are multiple dimensions, there is a "space" in which things exist: pitch space, timbre space, harmonic space. This space is an important part of the context; making the space obvious is therefore an important part of any presentation of those entities.

Structural vs. statistical approaches to description. One way of describing "what happens" is functional: B happens whenever A happens. Another way is statistical: A happens 45% of the time, B happens 65% of the time. To get a complete picture, you need both of these.

Identify formal vs. expressive aspects. One of the shortcomings of the conventional approach to "music theory" is that it deals with the anatomy of music more than the effect. A complete theory has to have both.

Different (distances of) views of complexity to see different levels of detail. A basic principle: you need to see the data from lots of different perspectives. In the Music Animation Machine (MAM), for example, the close-up view shows details of the performance (a rolled chord, say), and a far-back view shows the long-term (say, harmonic) structure.

Data as tool. The first time I was made aware of this was in a lecture by Rich Gold: he had a computer program in which a mathematical "surface" of varying heights was designed by the user, and then a rotating pointer was moved over this surface; the changes in height were translated into the waveform. Another example is a paint program in which a brush or other tool is drawn by the user, and then in turn used to draw other things.

WatchMe/Interactive/DoMe. This spectrum of interaction was also shown to me by Rich Gold. In music, a concert, is Watch Me. Jazz in an ensemble is Interactive, and a solo performance is Do Me. In designing any "interactive" device, it's worth considering expanding the range of interaction to cover this entire spectrum, so that a person can interact to the extent they please, and feel neither overwhelmed by too great a sense of responsibility (as in the violin) or bored by too trivial an involvement (as with the record player).

Is it possible to define the atomic action/element? If there isn't an obvious answer to this, then it's potentially a useful question to ask. As with many questions, there are often many answers.

Vagueness and graininess. When we look at a photograph up close, we sometimes see the grain of the paper. We wouldn't say "that's a picture of something vague," since we understand that things are not vague — only representations can be vague. When presented with a description, we sometimes say "that's vague," and wouldn't tend to say "that description is grainy." But is there such a thing as a grainy thought?

Can a map be made? Often a useful exercise.

What are the economic considerations? This is not to say that the economic considerations should dictate or even merely govern. But an economic consideration is a constraint like any other, and constraints, even ones that are only considered temporarily, can focus the inquiry.

Can a proof be made? It's handy to have a proof in your pocket. It may not convince anybody, but it might get their attention (which could be more useful).

How can it be notated? Again, just as an exercise, developing a notation is a good way to become familiar with the features and structure of the thing.

What are the various time-scales at which things happen? Often, one time-scale tends to dominate our thinking. Even when this is appropriate, there are usually things at other scales which are worth considering.

How can we distinguish between purposefulness and coincidence (numerology)? There are only a limited number of "small numbers," and it's unavoidable for them to come up in amusing combinations. Sometimes, these combinations shed light on some underlying structure; sometimes not. When you come across a coincidence, it's better to find out what it means (or demonstrate that it means nothing) than to merely remain in awe.

"News" vs. "history." In perception, "news" is sensation and "history" is memory. What's the right way to think about perception, then: as news or as history? Is our perception of a piece of music a series of sensations, or is it the memory of those sensations? Both, clearly. Does our description of aural events mirror this duality?

How much to distort the interface to match the user and how much the reverse? The shape of the interface can give the user information about the shape of what lies beyond the interface, which is good. Or, the interface can be fitted to be comfortable for the user. Ideally, both of these can be done, but that is not always possible.

The software spectrum. At one end of the spectrum is machine language: you tell the computer exactly what you want it to do. At the other end is a programmer: you tell the programmer exactly what to do. At the machine language end, you have complete control, but it's too much work. At the programmer end, you are stuck with whatever the programmer implements. To cover the middle ground we have "high-level" programming languages (which simplify the task of programming) and multi-purpose software (which make programming unnecessary). Still, though, programming is harder and software less flexible than they should be, and it's worth returning to the question of how to make it easier to make software that does what you want.

Metaphysical bug tracking. Metaphysics, especially the personal metaphysics every person is involved in, asks "what is reality?" There are parallels between this activity and the activity of tracking and fixing bugs in software. In both, there are points where you run into anomalies (apparent contradictions). In bug tracking, there is a formal procedure for reporting anomalies (contradictions among product specification, product behavior, user expectation, etc.). Could this procedure be made more general, such that it could applied to any endeavor? A related generalized model is the item/responsibility/disposition model of project management.

1.2 Motivation, Pedagogical Psychology

It's *fun* to know what's going on. We thrive on novelty, but novelty can only be understand within a context of familiarity. To present an unfamiliar thing effectively, then, you want to make sure that everything besides the new thing is familiar.

What happens when you don't know what's going on? There are three possible responses. You can **ignore** what's unfamiliar. You can **store** it (make a note that there's something you don't understand, with the intent of returning to it later). Or you can **explore** it. In teaching, it's best if the student can respond to the unfamiliar by exploring. Second best is storing; worst (as it leads to disorientation, frustration, and loss of motivation) is ignoring.

A good principle + a good story = a great teaching piece. Without the story, a good principle is something that, in hindsight, you're glad you've learned, but a teacher needs to understand how to tempt the student forward.

How to build a stimulating learning environment (esp. with/inside a computer). In social situations, the intellectual rewards of learning can be supplemented with interpersonal rewards (e.g. a teacher's positive reaction). Is there any adequate equivalent for software-directed learning? It's worth considering why people (especially children) like to play computer games. These require some learning, but there is a large amount of practice. What distinguishes the "practice" of playing video games with the "practice" of working on scales on the piano — or multiplication tables?

What removes enthusiasm/curiosity? Every teacher has had the experience of realizing that the student has lost interest in the presentation. Painful as such occurrences are, they are valuable feedback.

The function of desire in learning. It is possible to learn without desiring to (and not learn in spite of desiring to), but it's generally assumed that desire (and things related to it: curiosity, the need to feel in control, to expand one's sense of self-worth) play a part. How can a student's desire to know be provoked? In some fields (e.g. mathematics, physics, etc.) it's possible to provide a puzzle or anomaly which focuses a student's will to know; are there such provocations in other fields?

How can you learn how much improvement is possible? If people knew how much they could learn (and how much they would benefit from learning), they would be more motivated. Knowing that it's possible to learn something is an important step in the direction of learning. How can you demonstrate to a student that they are able to learn?

Autonomy. People are healthier and happier if they can make their own decisions and do things on their own; autonomy is an important principle in learning, since learning is the process of making something your own. The Taking Children Seriously movement takes this principle seriously; the child as an autonomous thinker is something they very much intend to foster.

Distractions. There are many kinds of distractions. In the lesson itself, there can be distracting elements. The teacher's personality can be distracting. The student may be thinking about other things unrelated to the lesson., or may have psychological problems (fears not directly caused by the lesson). There may be noises or objects in the environment. In some cases, the distraction can usefully be made into a legitimate focus of attention.

Rote vs. cognitive vs. creative skills. We teach children to count by doing it with them, over and over, and so it's tempting to think of counting as a rote skill. However, recognizing that something can be counted is a cognitive skill, and deciding *to* count something (deciding that it's worthwhile) is a creative decision. Many skills have this range of rote versus creative application, and it's not appropriate to think of any skill as purely rote.

Collecting. Learning is usually characterized as a mixture remembering and skill acquisition. However, when we value what we're learning, the objects of our learning become more tangible; a juicy piece of gossip has an almost tangible quality, as does a long sought-after answer to a puzzle. Just as some people are packrats, some people recognize the possible future value of ideas. How can this tendency be promoted in students and in oneself?

1.3 Specific Ideas, Questions, Problems (but not quite projects)

1.3.1 Music

1.3.1.1 Rhythm

How to show rhythm, how to represent it meaningfully. There are several common methods for notating rhythm: symbolic (in which the notation represents rhythmic concepts which must be learned ahead of time), graphical/linear (in which events are plotted in time), graphical/differential (in which rhythmic deviations are plotted), etc. However, the difficulty people have in learning to interpret these forms of notation suggests that there is progress to be made on this front.

Use a circle to represent rhythms which repeat. This is simple and obvious, but it works surprisingly poorly. The problem with constant circular motions is that the viewer's sense of the impending arrival of an event is not very strong: it's possible to follow circular motion and anticipate collisions on a circular path, but it takes a lot of attention. A possible direction of research, then, would be: how can we visually indicate "it's coming!" better. The techniques of conductors (who never use constant circular motion!) would be worth studying.

Use size/time metaphor, shrink/grow as period changes. A simple way to do this would be to use smaller and larger circular paths (though see previous item); you would zoom in to see how a rhythm worked in a sub-phrase, and zoom out to see how a form worked.

Wave motion could be used as analogue of regular beats. Wave motion (e.g. in water) is more predictable than circular motion; the crests and troughs are more distinct. Could waves be made more complex, sophisticated, or subtle while still being obvious and comprehensible?

Does one beat predict/suggest the next, or reflect/echo the previous? Each time a rhythmic pattern repeats, it confirm our sense that it *has* repeated, and strengthens our anticipation that it will repeat again. The two common ways to represent aural rhythm visually are to duplicate it (i.e.: to provide visual events in sync with the aural events) or to convert time into space (as in bar-graph notation or conventional notation). However, both of these leave the confirmatory and anticipatory qualities up to the viewer. Could these be depicted visually?

The physicality of music making/response. Listening to music is more than just cerebral; it may remind us of (and even give us of sensations related to) physical activities such as dancing, walking, running, swimming, massaging, drawing, having sex, eating, etc.; music making itself is even more suggestive of such activities. Conversely, these activities can be used as a source of things which are suggestive of music.

1.3.1.2 Pitch, harmonics, timbre

Analogies to our pitch sense. The easiest approach is to consider pitch as a simple, linear quantity. However, this ignores many aspects of pitch. A line continues indefinitely, but our pitch sense does not; at the low end, pitch is harder and harder to discern, and finally turns into vibration and rhythm; at the high end, pitch also becomes hard to discern (and there's an upper limit). A sound which we think of as having "a pitch" usually has many frequency components, leading in some circumstances to ambiguity of what pitch is present. Points on a line are not precisely memorable, but we have a strong pitch memory; some people can remember pitches forever (perfect pitch), but most people can remember pitches for a long time (which makes tonality possible). There is the phenomenon of octave equivalence. A sequence of pitches is memorable in a way that a sequence of positions on a line is not. Each of these ways in which pitch is more complex than a simple linear quantity suggests possibilities for ways of analogizing pitch.

The problem of 7- (or even 5-)limit when using color for pitch "class." There is only one aspect of color for which we have anything close to a reliable memory, namely, hue (i.e. angle on the color wheel); other aspects (saturation, brightness) are much less memorable, and are thus less suitable for symbolic use. Unfortunately, hue is a linear quantity, and while pitch class is linear in the frequency sense (i.e. the chromatic scale), there are several important measures of pitch "closeness" other than the chromatic half-step: the perfect fifth, the major third (and perhaps others). Given the simplifying context of 12-pitch equal temperament, all pitch relationships are circular (e.g. three major thirds returns you to the starting pitch); the challenge is to depict all the important ones, as circular, simultaneously.

The integrity of a musical "object" (fundamental + harmonics + other). An important task of auditory scene analysis is the assignment of harmonic energy at various frequencies to the appropriate sound objects. We do not hear all the harmonics of a tone separately, but as a fundamental which "owns" a bunch of spectral components. Standard visualizations of sound take one of two approaches to this complexity, either (as in the MAM) show only the fundamental as a single sound object (stripped of its spectral components), or (as in a spectrogram) show all the harmonic energy, without providing any means for associating a fundamental with its spectral components. The Voice Tracker display associates the harmonics with the fundamental, but it only does it for a single sound at a time, and the approach can not be extended to multiple instruments in a useful way. What is desirable is a way of showing the harmonics both as distinct (so you can see when harmonics are shared by multiple sound objects) and "owned" (so that you can see the timbre of each sound object). Related to this is the idea of having each frequency component "seek out" harmonics related to it, which it would then bid to own or be adopted by.

Affects of the harmonics. It has been mentioned that various classes of harmonics have different affects: 2^n harmonics sound relatively transparent; $3*(2^n)$ harmonics sound strong; $5*(2^n)$ harmonics sound warm or sweet, $7*(2^n)$ harmonics sound edgy, etc. These distinctions could be incorporated into a visualization of timbre.

1.3.1.3 Interpretation

Compared to studying the compositional structure of music, studying musical interpretation involves a greater attention to (often hard-to-define) subtleties such as differences in the timings of note onsets and releases, dynamics, etc. Tools for studying and editing musical interpretation must make subtle differences visible, but that is not enough; they must go beyond that and present the differences in such a way as to show why differences *make* a difference. For example: at the end of a piece of music, some ritardandos feel right and some don't, and it's been demonstrated that those that feel right are of the same form as the ritardandos characteristic of human locomotion; however, it's next to impossible to judge, by looking at a series of spaced lines representing note onsets, whether a ritardando being depicted follows the human model. How could this be remedied? Likewise, it's difficult, looking at a bar-graph score, to see what's going on in the articulation of a phrase, or in the placement and intensity of accents (amplitude or agogic).

In addition to showing the subtleties involved in interpretation, a tool for the study of interpretation also needs to show the structural details that the interpretive details are meant to delineate. For example, an accent may be used to bring out a change in harmony; the tool needs to show the harmony and the accent in order to make relationship between them obvious. Likewise for showing structural details that are delineated by phrasing. Ideally, this approach would lead toward making the *feeling* of music visible, and would let us understand, more directly, what contributes to the most satisfying, most convincing, most *right* rendition of a musical gesture (something which is currently so tangible and yet so elusive).

In a performance editing tool, it's not sufficient to provide control of the timing and dynamic level of each note in the performance; without higher-level editing operations, all but the most trivial changes involve a burdensome number of notes and attributes of notes. The master class is a useful metaphor for how changes could be made: instead of making individual changes, you specify the nature of the change, or give an example of how something is done, and have the relevant details incorporated into the rendition. In a stripped-down implementation, you could select notes, and then re-play them, and have only certain selected aspects of your new performance (pitches, accents, articulation, etc.) applied to the rendition. The conductor program is another approach to this (and see **Less Work** in the **New Ideas** folder).

1.3.1.4 Notation, Scores, Sight-reading

Evaluation of conventional music notation. It is certainly true that, in a Darwinian sense, conventional music notation has proved itself to be "the fittest." However, little work has been done to define, precisely, the requirements a music notation must satisfy or measure the extent to which conventional notation satisfies them. Developers of alternate musical notations are hindered by this incomplete view, and their results are predictably unsatisfactory. Part of the problem is that musical scores serve many purposes. A score serves as a record of a composer's compositional and interpretive ideas; it is also used by performers to learn the piece from, and to perform from at sight; it is also used by students of music theory, and by listeners. Each of these users of music notation has different requirements. The MAM is for listeners (which could include composers, performers, and other students of music—but *as listeners*), and thus emphasizes different things than conventional notation (which grew primarily out of the needs of composers and performers). Still, the MAM was developed in an ad hoc fashion, without a clear understanding of what the needs of listeners are; it would be good to identify those needs and develop new notations from first principles.

Animating musical scores. It is generally agreed that, in some sense, music "moves," but little agreement about what this means—or what it ought to mean to "animate a musical score." Could a musical score be self-animating, that is, could the information in the score provide everything that's necessary to generate a satisfying animated analogy to the music? There have been many attempts at animating scores, but the most successful of these (e.g. Lejf Marcussen's *Tonespor*) have been shaped by ideas of the animator, and not generated completely automatically.

Sight-reading. What does it mean "to sight-read" music? What does a performer need to know to be able to sight-read fluently? Do all competent sight-readers do it the same way? How can it be taught? The experience of sight-reading music fluently is among the most intense—and intensely satisfying—available in music, as it combines the challenge, stimulation and aesthetic reward provided by novelty with the pleasure of anticipation and familiarity (you have to be familiar with music and anticipate where it's going to go to be able to sight-read it); how could this experience be extended to non-musicians?

1.3.1.5 Instruments

A musical instrument's interface works both ways. The study of musical instrument interfaces usually focuses on what they let the performer control. However, an interface works in the opposite direction, too: to tell the performer about the control it offers (and, by extension, about the musical elements which are being controlled). For example, a flute shows you a diatonic scale, a piano shows you the relationship between a diatonic scale and a chromatic scale, and a steel drum shows you the circle of fifths. This idea can be extended to instruments specifically designed to teach individual musical lessons (and, especially, to hyper/virtual instruments, which are less constrained in their designs and often less expensive to make). [Section on Bill Collins' devices goes here if he gives permission.] An important aspect of a music instrument's interface is how it *feels*; it is different to play timpani than it is to press keys on a keyboard, even if the resulting sound is the same. Also, the shape of an instrument is important; ideally, one's position on an instrument's interface is discernable by touch alone; this is a weakness of interfaces that are based on computer screens (e.g. touch screen); wind instruments don't require any looking; piano and guitar require relatively little, harp requires a lot.

Computer-based musical interfaces. There is a lot of work to be done here!

The human (inter)*face.* We use some small parts of our face (our lips) to control musical instruments, but we have exquisite control of lots of other parts of our face. With sufficiently sensitive optical recognition, these other parts could also be used to control aspects of a musical performance.

Sharing musical responsibility. The most common "two-person musical instrument" is simply two instruments, but there are many other ways of sharing responsibility. Conducting students are trained by having one student conduct while another performs. With a computer mediator, responsibility could be shared in other ways; you could, for example, have one person for each of the "local control" aspects of a performance (rhythm, loudness, melodic shape, harmony, texture, timbre, instrumentation), with another person to direct (control overall form).

Convenient access to axes of control. The difference between one instrument and another often lies in differences not in what can be controlled, but in how and how easy it is to move through the range of control. For example, a recorder is like an organ pipe, but to change the tuning of an organ pipe takes several minutes, and can therefore not be used expressively. As we move toward more and more synthetic generation of musical sounds, we are often putting control of those sounds outside of the real-time reach of a performer. In 2001jul Keyboard magazine (see **Interfaces** file), page 20, the "Keyboard of the Month" features an add-on to a synthesizer which brings out-of-reach controls closer to the performer. How could computer interfaces be moved in this direction?

1.3.1.6 Visualization

The purpose of music visualization. There are many possible applications of music visualization. My goal is to make music visualization for which the question "what is it for?" is as unlikely to be asked (and its answer as obvious) as it is for music or food or sex ... That being said, some areas of application are worth special consideration. Going from the more obvious to the less: the MAM score can be used a substitute for conventional notation, or as a way to study details of an interpretation; viewed by pre-literate children, the MAM score is a window into the structure of music (a preparation for music notation); music visualizations can be designed to show a particular compositional aspect of music (scales, harmonic motion, tonality, texture, instrumentation, rhythm, etc.) to help a student focus on that aspect; likewise for a performer, a visualization could help them focus on certain aspects of their performance (tempo, phrasing, dynamics, etc.). In the context of music production, music visualization is a necessary part the user interface for editing and mixing operations.

Fundamental research. Superficially, music visualization would appear to be a creative, as opposed to a research, activity. However, vision is very different from hearing, and the differences between these senses play a big part in determining how music can be portrayed visually. For example, a visual object, even a moving one, has a form that exists instantaneously, independently of time, while a sound object (though it may have attributes that are perceptually instantaneous, such as pitch or timbre) cannot be said to have a form that exists in an instant; therefore, a visual analogue of music will necessarily have forms without precise parallels in the music itself; there will always be some asymmetry between the two. There are many other mismatches between vision and hearing (the speed of responsiveness, the effects of combining components, the structure of the dimensionality, etc.), each of which requires design decisions and compromises.

The kinesthetic quality of music. Music can make us want to move, and it can make us feel as if we were moving. It can make our muscles feel tense or relaxed. Both hearing and vision can impart the sense of motion, either literally (e.g. a sound panning across a stereo field, an object moving) or figuratively (e.g. an ascending scale, or a picture of a zigzag). In a music visualization, it's useful to ask "what is moving in this music?" and search for ways to express that motion. Motion can be examined at various time scales: the shape of a single note (its envelope), the motion of moving from one note to another (a "step"), the motion that takes you through a motif (a "gesture" — see MAMgesture.doc) or melody (a "path"), the motion between sections (form). The sensation of motion can be engendered through melodic, harmonic, or rhythmic means.

Visualizing musical constraints. Within any style of music, certain things are more likely than others; e.g., in atonal music, combinations of notes occur which almost never occur in tonal music. Using the MAM's "harmonic coloring" mode, it's possible to see whether a piece is tonal or atonal. Visualizations could be designed to show a piece's conformance to (or deviation from) other constraints (e.g. triadic harmony, staying within a diatonic context, canons, rhythmical regularity, ranges of instruments, inter-instrument intervals, melodic motion, etc.).

Specifying the "note." In designing an animation, there are various ways to think about "a note." Simplest is the fact of a note during the time it is sounding. There is the melodic lead-up to the note. There is the performer's preparation for the note (and the listener's anticipation); likewise, the follow-through (and memory). There is the note's position within a chord. One way of thinking of a note is as a position in a path that an instrument follows (or as a segment in a series of actions done by a performer). Another way is as a creature that is alive all the way through the piece (though specially active during and around the time it is sounding). Another is as a pitch, which is active at various times through the piece. Another is as a position within a melody (which may get played several times during the piece). In the context of a visualization created from an audio signal (rather than performance data), there is also the issue of identifying a note within (and isolating it from) its context.

Synesthesia. See MAMsynesthesia.doc, Synesthesia in Pendaflex files.

1.3.1.7 Education

What is it that musicians learn? What is "musical knowledge"? A major difficulty in becoming a musician is defining what it is that must be learned. When people learn more than one instrument, they usually find the second instrument easier to learn than the first, the third easier than the second, etc.; this suggests that there is something non-specific that's being learned (and which therefore does not have to be learned over for each instrument). Likewise, practice in musical composition helps with interpretation in performance. What is this "under-knowledge," lying beneath what is ostensibly learned? Could it be taught separately? Could a person in this way become "musical" without learning anything specific to a particular instrument or musical activity?

What musical knowledge is essential and what is inessential? Presumably knowing the fingering of scales on the piano is less essential than being able to hear music in one's "inner ear" before one plays it. What is inessential, we should relegate to a machine; what is essential, we should learn (perhaps with the help of a machine).

A beginning music student has special needs. Music is *complicated* and cannot be learned all at once; it is also *difficult*, and requires practice to become adequately familiar with each part. Therefore, a teacher (or teaching system) must break the learning into small enough units, and encourage the student to spend enough time working on these units to attain mastery.

1.3.1.8 Legal

Copyright of musical performance. What is a copyrightable musical performance? Recently, the music industry has begun to refine this question by asking, in effect, "what is the 'atom' of copyrightable musical performance?" and focusing on the performance as an *audio signal,* the answer being expressed as some number seconds. However, there are other axes of atomicity than time. What happens if you transform an audio signal of a performance into a MIDI file? (MIDI stands for Musical Instrument Digital Interface, a standard for the exchange of musical data between instruments and computers.) Does that file violate the copyright? What if, after doing this transformation, you play it back through an instrument that sounds completely different (or manipulate it algorithmically), so that the "average listener" would be unable to recognize it as being related to the original? In the long run, we will learn that the main value a performer brings to a performance is not the recognizability of particular clips, but an aesthetic. What happens when we are able to extract and reproduce a performer's aesthetic? Can a person's rights to their own style be protected by copyright? (If the precedent established by DNA patents is any indication, probably not.)

Licensing CD recordings for use in music visualization videos. This has long been a stumbling block in the path of making better MAM recordings. Recently, services have arisen which will sell anybody a license to various "generic" recordings of classical music.

1.3.2 Language

Meaning. Chemistry was the first language: an organism would communicate its "essence" by releasing chemicals that were distinctive to it. This is an expensive way to communicate, since you must expend matter. In "modern" languages, you only have to expend energy—which Einstein showed is equivalent to matter, but it's cheaper. A small physical gesture can suggest a larger action (a shaking fist can mean "I will not hesitate to beat you up"). As language becomes more abstract, the gesture resembles less and less the thing it stands for. But, fundamentally, does language ever develop meaning which is differentiable from gesture?

Inflection. Inflection provides a means of generating (and organizing) variants on a root word; this technique is flexible (every root can be inflected in many directions) and powerful (when you learn a new root, its scope of application is multiplied by the number of possible inflections; when you learn a new inflection, its value if multiplied by the number of roots). Inflection can be provided by gender, number, declension, conjugation, pre-/in-/suf-fixing, or by idiomatic combinations (look up, look out, look down, etc.). The idea of inflection can be applied in many areas; for example, the keys on a typewriter keyboard are inflectable (via shift, control, and alt keys).

Phonetics. In spoken language, "motion" is the motion within and between phonemes. As in music, there are points of tension and relaxation, various types of articulation, more and less natural transitions.

1.3.3 Sensation

What are the dimensions of the senses? All senses involve time (since all sensations can change over time) and amplitude (proprioception being a possible exception); beyond that, things get more complicated. Analogies between senses are based on mappings; for an analogy to be practical, the dimensionalities of what's being mapped must correspond. For example, an analogy between smell and musical pitch is impractical, but an analogy between touch and vision is practical. (The more fundamental question "what is an analogy?" is still open, however.)

One way to consider the dimensions of the senses is to define attributes. Hue, for example, is a visual attribute. An attribute may involve more than one dimension, and so may not be completely orthogonal to other attributes, but there are contexts in which practicality is preferable to precision.

An interesting thing sometimes happens when we study of our sense organs. As part of answering the question "what information is coming into this organ, and what happens to it?" we may make a model and then observe it in action. In doing this, we've converted one sense into another; we may, for example, see the activity of the cochlea (or the cells which process the nerve impulses originating at the cochlea). Or we may hear the activity of nerve impulses coming from the optic nerve. So, we have a kind of artificially induced synesthesia.

1.3.4 Computer

Training the user interface. Ideally, a computer does exactly what you want, as quickly and easily as you can conceive of it. At its most fluent, typing can be like that, as can computer drawing and painting. Few other things come close. The spectrum between designing software (writing code) and using it (limiting yourself to operations conceived by the designer) is worth exploring. Macros are somewhere in the middle: you tell the computer "this is something I want to do; remember it." There are problems with macros, however: to make them flexible requires an effort which is more toward the programming end of the spectrum; when you create lots of macros you have trouble remembering them; etc. These are not fundamental limitations, but rather opportunities for creative design; for example, macros could be collected continuously, parsed and organized automatically, and presented to the user in a way that reflected the structure of the macros, their frequency of usage, etc.

Point and select can be simultaneous. In the MAM editor, Select (an operation) and Point (to a place to apply the operation) are done by different hands, and can thus be done almost simultaneously. For high-performance ("power") user interfaces, this approach can speed up user interaction by more than a factor of two.

Combining features. Often new classes of software arise from combining features of existing pieces of software. For example, the spreadsheet (the first "killer app" for personal computers) combined a text editor (with support for columns), a calculator, and a programming language.

"Interactive" software. It's a truism that good software is interactive, but there are no well-established metrics for interactiveness, so it's hard say whether a user interface is optimally interactive. Still, there are various rules of thumb. Interactive means "you can see what you're doing." There are degrees of this. When people were typing on teletype terminals, they thought they could see what they were doing because the letters they pressed on the keyboard were typed onto a piece of paper in the machine. But with the "glass TTY" (teletype), found we could see how the paragraph was changing as we edited it, and with DTP (desk-top publishing), we found we could see, with ever greater precision, what the final printout would look like. But there is still a long way to go; we cannot see what our words mean, or the effect they will have on our readers; for those things, we still have to use our imagination. So, there's still room for improvement.

A related question: how can a display become a tool? A picture is useful; the more interactive a picture is, the more tool-like it is. A picture can be made interactive by allowing the user to: change their point of view (or point in time which is depicted), change the data that generates the picture, add or remove features from the picture, etc.

Computer-based learning environment. We do not yet know how to make good ones. We have computer-based environments that are stimulating, useful, comprehensive, rewarding, interesting, fun, demanding ... but rarely all at once. For some very limited kinds of skills (e.g. typing), there are adequate solutions.

1.3.5 Mathematics, physics

Making statistics (and mathematics generally) tangible. This is the next frontier of mathematical visualization (now that we've done geometry, algebra, set theory, calculus, etc.) and, given how we are increasingly having to make decisions based on contexts so complex that the only possible way of dealing with them is statistically, perhaps the most important.

And it's a good idea, generally, to make mathematics more tangible. It's been proposed that African-Americans are less interested in computers than members of other ethnic groups are because computers aren't "black enough" (meaning they're too nerdy, and not physical enough, not cool enough); the same could be said of mathematics. How is it that music and language can be cool but the language of mathematics is not cool?

Dance can be viewed mathematically: it involves number (especially dancing to music), motion, geometry, etc. How could dance be made more mathematical? In Dalcroze Eurhythmics, musical training is combined with dance; why not mathematics?

The most common way to describe motion mathematically is by specifying velocity and direction. While this may be the simplest, it is not always the most humanly meaningful, because we don't have a "velocity organ." Rather, we have muscles, activated by nerves (initiated by intention), pushing on weights arrayed on levers. How do we use this system to move from point A to point B in time T? It is much more complicated than calculating the velocity (i.e. distance(A, B) / T)—and much more interesting!

Limits of DSP. There's a fundamental relationship between sampling rate, settling period, frequency, and resolution. You can't get information instantaneously!

1.3.6 Other

Content vs. View. Content is usually considered to be distinct from our (point of) view of the content, but there are ways in which this distinction breaks down. If we consider "what we're looking at" to be "everything," then point of view totally determines content. If "point of view" is taken to be "conditions under which the object exists," then point of view can likewise change that object significantly.

A sentence is a solution to a topological problem. Language is linear, but what it describes is not linear. Similarly, although the perception of language happens sequentially, what's constructed in the mind of the listener is not linear. The process of communication via language is like building a physical object: there are certain ways of building that work better than others (e.g. if you're building an arch, you can't start at the top). Designing a piece of writing to build a non-linear idea is thus a solution to a kind of topological problem (see **Topology in Writing**).

Real-time decision making. We humans are real-time creatures, but we are not very adept at learning real-time behaviors, especially real-time decision making. When we practice video games, for example, we do not practice how to make complex decisions in real time; rather, we learn how to see our situation in a more sophisticated way, so that we *don't have* to make decisions consciously; instead, we "just see" what we need to do. What kind of training environment would it take to make conscious decisions more quickly? (see notes in **New Ideas** folder)

Gesture is multi-modal. For example: you can point *at something*, you can point *with emotion/motion*, you can indicate a certain *distance* with a gesture, etc.

One-handed keyboard. Sometimes when I'm typing words that happen to have letters for only the left hand (see **LeftHandWords.doc**), I wonder whether a useful one-handed keyboard could be built? Some languages have fewer letters than English. If we approached it phonetically, we might be able to do with fewer. We could develop a dialect of English that didn't use UJMIKOLP (and have the left hand home keys be SDFG instead of ASDF), and type without punctuation (the computer could automatically translate between our dialect and true English). Etc. There are so many situations where a one-handed keyboard would be useful!

Diagramming tense. Tense in language has to do with the relative position in time of the speaker and certain other events; the positions of tense could be shown in diagrams; such diagrams would be useful when learning a language (especially in cases where tense in the student's native language and the language being learned do not map one-to-one), and could be part of a completely graphical language.

Gesture. Has anybody developed a comprehensive, formal method of specifying gesture?

The Aesthetics of Humor. A joke is rarely as funny the second time (except for young children—which raises the question: do they learn that a joke isn't funny a second time from others, or do they develop this on their own?). This fact seems to be related to other aesthetic issues: the nature of surprise, repetition.

Driving. Like language, driving is a complex activity that many people are expert at and can do without conscious attention to the details.

Biofeedback. It's probably a little early yet for useful biofeedback. (see **Biofeedback** pendaflex file)

2. Projects

2.1 Music Related Inventions, Software, Other R&D

2.1.1 Music Animation Machine (MAM)

2.1.1.1 MAM hardware/software features/development

The MAM itself. I've completed two versions of the MAM software (MAM1-1985, MAM2-1988), started a third (MAM3-1995), and contemplated others (see **MAM2000.doc** and **MAM4.doc**). I considered writing a specification for an application framework and hiring Russian programmers to implement it (see **MAMVIEW.doc** and **MAMViewReqCapture.doc**) as a way to overcome my inertia (and lack of time), but I couldn't convince myself that their references panned out, and while I didn't mind getting scammed personally, I didn't want to encourage them, so I dropped it. I can't decide what should be in it, who should have it, etc. Should I try to make a business of it? Should it be a freeware performance editor? At this writing, I'm considering making some new videos, and it would be nice to have better software for that project.

Synchronizing the MAM with CD recordings. A shortcoming of the MAM to date is the relatively amateurish quality of the performances; one way to sidestep this would be to use commercially-available recordings with software that could synchronize MAM displays to them (see **MAMSYNC.doc**, **MAMsynchronizer.doc**, **COMPACT DISK** in the pendaflex files; **V09** is where I started to work on this). I contacted various recording companies to ask how they'd feel if I made "companion" software to go with their CDs; the answer was basically "we forbid you to do that."

MAM patent. Although the MAM cannot be patented (because it's been public for too long), it's useful to consider what a patent for it would look like (see **MAMpatent.doc**), and to think about variants that could be patented.

MIDI/MAM Gamelan. The idea here is twofold: make a MIDI-fied gamelan (I made some prototype MIDI mallets which worked okay), and develop a form of MAM notation for studying gamelan performance practice.

TAPPER as part of the MAM. Currently, there's a mode in which the MAM does the Conductor Program. This could be incorporated into the performance editor, too, with such variants as Apply Rhythm, Apply Phrasing, Apply Dynamics, etc.

Interactive MAM. The MAM could serve as a performance interface: the score would scroll by, and you'd interact with the notes to shape them. The easiest way to approach this would be with a touch screen, but better would be if the experience were 3D and full-body.

Fanciful animation, imagery. This of course is wide open.

Performance editor improvements. The MAM editor has some nice features, but compared to modern sequencers, it's missing a lot, too (see **MAMeditorComparison.doc**); it would be good to modernize it.

MIDI-to-MAM converter. The MAM currently uses a proprietary file format. Ideally, it should be converted to use standard MIDI files (though there are implementation issues with this); as an interim measure, a file converter could be written.

MIDI-to-MAMQuickTime converter. This would convert a standard MIDI file into a MAM-style QuickTime movie (like the In Dulci Jubilo QuickTime movie that's now online). The user interface for this would allow the user to adjust the display various ways, give a preview of the movie, etc.

Fill empty background. The MAM background is largely empty space; this could be filled with something (e.g. video of a performance).

Non-note MIDI data. Currently, the MAM displays pitch and timing; it could be modified to include other MIDI information (e.g. key/channel pressure, control, pitch-bend, etc.). Another approach to this would be to have an event-list editor.

Wall-sized display. The few times the MAM has been displayed on a larger than normal screen, the effect has been good. The MAM display is fairly simple, so it scales up well. This could be done with David McCutchen's immersive video display. Alternately, a display on a long wall (as in Union Pacific "Star Wars" room, see picture in files) is a nice way to deal with long time intervals (like the original MAM scroll).

Measures. Currently, the MAM display makes no reference to measures; adding them would be useful for navigational purposes, and necessary if the software is to do editing algorithmically.

Cut and paste. Duh.

External control. There are times when it's useful to sync a MAM recording to external events. Currently, it will sync to MIDI; this could be expanded to include syncing to audio, network, etc.

Channel/track support. The current MAM supports a crude MIDI channel redirection; it would be good to have a more flexible model, in which attributes were assigned to tracks (rather than MIDI channels). Attributes would include: color, style of display, input channel(s), output channel(s), visibility in the display, etc.

Insert note. The method for this is currently clumsy.

Visual feedback. There are several things that are difficult to see or invisible: cursor, being in/out of edit mode, what edit/play mode is active, etc.

Display ordering. Since notes are layered, one channel's notes may obscure another's. Two features would help with this: channel display ordering (decides which is frontmost, next, etc.), and channel display hiding.

Keyboard display. To help with note identification, a keyboard graphic could be positioned at the left edge, or slid. Alternatively, the display itself could be "staff-ized" (temporarily draw pitch lines and accidentals).

Filters. You'd select a range of notes and apply an editing operation.

Print/snapshot. You'd select a range, and dump it to disk as a .bmp file.

Channel/track names. Typically, you'd assign names relevant to the score (e.g. right hand, bassoon, percussion), or to playback (e.g. PROTEUS 35 PAN 2 VOL. 50).

Multi-window display. This could be used to compare two performances, or to assign different "spaces" to different instruments (where direct overlay is inappropriate, e.g. for gamelan).

Note shape. I've done some simple experiments with this (CONTEXT0: draw and age stars, CONTEXT1: stars, squares, circles, rhombi, CONTEXT5: draw and age circles but don't move them, CONTEXT6: good-n-plenty-shaped notes) in a stand-alone context; this could be supported on a by-track basis. (And these could use improvements in and of themselves.)

Speed control. This would be used to synchronize the MAM display to a live performance. The "Cranker" interface would be a good one (sensitive).

Non-linear scale. In the original MAM, time and pitch are linear. The "now" can be emphasized by giving recent/impending time more horizontal space (see **V25**) or more vertical space (see **Fisheye MAM**). This could be applied to other kinds of display than the standard MAM.

Paper scrolls. The original MAM scroll (see) had aspects that were lost in going to a VDT display. It would be nice to have a way of producing long paper scrolls.

2.1.1.2 MAM tapes

More MAM videos. I've planned to do more videos ever since the First Demonstration Reel. Contemplated projects include: more general-interest collections, collections specifically for use in music history and music appreciation classes, counterpoint (especially Bach fugues), Chopin etudes, jazz (see **MAMpieces.doc**), Godowsky, Debussy (Claire de Lune for David Miller, other things for Jamie at Classic Arts Showcase/CAS), the first Arabesque. In 1996 I did a tape as a commission, and in 2000 began selling it as the Second Demonstration Reel (which is advertised on the web and on a flyer included with the First Demo Reel). Also in 1996 I did a short documentary; but because it uses some copyrighted material, it can't be sold; a possible project would be to redo that, bring it up to date, and use copyright-free recordings (or my own).

Custom MAM tapes. At some point, I thought people might want this, and I put together a setup for cranking out MAM tapes (titling software with lots of font utilities, high-quality VCR, etc.), put together a catalog of pieces (see **list of pieces**), and put an ad in the Viewer's Guide that was distributed with the tapes I was selling. Several years went by with no takers, and I dismantled the setup. It could be a good idea, though; to do it right, I'd have to automate the production of tapes, improve the titling software, and market it right.

MAM production at home. Thanks to digital video, this is now well within reach.

Scan converter problem. One unsolved technical problem is that computer video scan rates are not quite the same as TV video scan rates. When you do "scan conversion" between the two formats, this discrepancy causes occasional dropped frames, which can be very noticeable if the MAM scrolling is fast. There are video cards which have TV format outputs, but there are compatibility problems (and other problems) related to them.

2.1.1.3 MAM web site

Improve layout. It's still hard to find your way around the site. It needs to be cleaner at the entry point (with thumbnails for the main topics), have a hierarchical index for the MAM stuff, and a flat index with everything.

Expand History page. This has been started, with nice thumbnails for all the inventions; this should be expanded to included both MAM and non-MAM. There could be a full-sized image of the scroll!

Add Bagatelle page. Art Weber proposed (see **Music for the Love of it** article) that Beethoven's opus 33 #5 Bagatelle is a musical depiction of a pinball game (pinball games were called "bagatelles" at one point); it would be nice to reproduce his article on line.

Expand the section on the conductor program, including other practitioners, implementation choices, video of a live performance (and write Rhonda Alexander when this is complete).

Put Malinowski scores (MIDI and PDF/Portable Data Format) online to download. There are a few compositions of mine which the passage of time has not yet ruined my opinion of; it might be nice to share these.

Add more downloadable MAM movies, and do things to let more people know about them; include shorter sections of pieces from the videos (and put thumbnails on the video ordering page); put a Keyspan ad on the final screen; include samples of CONTEXT-series and V-series experiments; track usage.

2.1.1.4 MAM applications, reflections

While it might be said that the MAM was first developed purely as **entertainment**, it is a kind of entertainment in which the pleasure comes largely from being able to recognize (and anticipate) things in the music which would otherwise go unnoticed (or unanticipated); it's a kind of entertainment that's related to cognition and learning. Many of its possible applications are related to this. It can be used to study **music theory**, musical **composition**, and

performance (see **ACaseForMamEditing.doc**). For very young students (including infants), it can serve as an **introduction to music notation**, and for older students, it can be **a step toward becoming musically literate**. Besides education and entertainment, there are other uses for the MAM. It can serve as a form of music notation for the musically illiterate; it could be used for sing-alongs, karaoke, etc. (see **Korean MAM Sing-Along**).

The **process** of developing the MAM (and software that followed) was interesting in its own right. A big part of this was a consequence of the MAM being unlike anything I'd previously done, and unlike anything anybody had done; this meant that I was making discoveries (and mistakes) completely on my own. As a result, some design problems (such as the editing paradigm) were solved in ways that still have advantages over their counterparts in contemporary music software, though as a whole, the MAM is primitive and incomplete (to quote Ashleigh Brilliant: "I may not be totally perfect, but parts of me are excellent.").

An unusually high proportion of those responding favorably to the MAM have been **architects**. Why is that? Is it that they are more sensitive to visual forms? Or does studying architecture make them more aware of form in music, and the MAM satisfies their desire to see what they've been hearing?

2.1.2 Musical interfaces

Freedom from acoustic considerations. Until recently (ca. 1000 years ago), designing musical instruments unavoidably required an artful balance of acoustic and ergonomic considerations (e.g. a flute could only play in a certain pitch range because its finger holes had to be spaced closely enough to allow them to be spanned by human fingers). The keyboard was the first big step in the direction of freeing instruments from acoustic constraints. Now, we are able to make instruments in with no acoustic constraints whatsoever: any interface can have transducers applied to it, and a computer can do any conceivable transformation of the gesture data it delivers. So far, however, we have not taken much advantage of this freedom; most musical instrument interfaces mimic preexisting instruments (e.g. MIDI keyboards and wind controllers). This is natural, since we haven't radically changed our concept of what it means to "play" a computer-assisted musical instrument. What new alternatives does a computer-assisted musical instrument offer?

The separation of musical concerns. In the first musical instrument, the voice, the performer was responsible for timbre (including words), volume, pitch, rhythm. The first musical instruments reduced the performer's responsibility for timbre and pitch. The first keyboard organ also reduced the player's responsibility for volume. In exchange for control of some aspects, the performer gets enhanced control in others: a keyboard player can play multiple notes at a time, and a violinist can change pitches much faster than a singer can. The question now is: what musical responsibilities are best left in the hands of the performer, and which can be usefully relegated to a computer?

Interfaces for the conductor program. In the conductor program, the computer is responsible for pitch, and perhaps other things. The question remains: what control to give to the performer, and through what interface? The follow possibilities are well-suited to application with the conductor program, but could be used in other musical instruments too. The **Lizard Trampoline** has a flexible surface which can detect continuous pressure in multiple locations as well as quick taps (though the use of conductive fingertips in the performer's gloves). **Rhythm Surface** is a clay-like surface which the performer both creates (by molding) and plays (by rubbing, tapping, squeezing, etc.) during the performance. **Ski Poles** are related to Stephen O'Hearn's controller design (see **Keyboard, 1999jan**, page 106), but uses two poles, anchored to the ground, and other **Body-sized musical interfaces** would allow a greater degree of audience participation (and perhaps greater control as well). **Bow Peep** and **Bouncer** combine aspects of bowing and drumming to allow both fast rhythmic control and envelope shaping. The **Music Harness** is an attempt to get more of the body involved in the interface (and to provide both position sensing and resistance). The **Hurdy-Gurdy Crank** is a way to tell a computer about tempo; needed improvements to the first prototype include: ability to "feel" notes approaching, dynamic control (perhaps by squeezing the handle), visual feedback (see **New Ideas** folder). **Tapper** (which I've implemented as both a stand-alone program and a mode of MAM2) is a simple implementation of the conductor program; improvements include: song position menu (so as to easily practice sections), support for trills, arpeggios, two-hand control, performance recording, performance modification of existing renditions, articulation interpolation. I should assemble some Tapper setups for various people who'd enjoy them (Craig Warren, Mervin Lane, Betsy Hanger).

Interfaces for beginners. Learning how to control a sophisticated musical interface is difficult; a second reason to delegate musical responsibility to a computer is to allow a beginner to focus on a certain aspect of performance so as to more easily master it. Infants, for example, have only crude muscle control, and are therefore not able to provide meaningful articulation with their fingers; an interface for them would use their larger arm motions, and perhaps limit pitches (or tonalities) to a manageable number while leaving them a full, subtle range of dynamic and timbral control. Tonality (in a 12-pitch system) is circular; an infant could sit in the middle of a circular, **surrounding tonality keyboard** to learn about tonal distance and direction. **Banger** (see **New Ideas** folder) is software that turns a standard MIDI keyboard into an instrument more suited to young children; this idea could be extended (with ideas from video games) to encourage non-mindless behavior (see **Creativity Enhancer** and **Anti-Stupidity Device** in the **New Ideas** folder).

Interfaces to teach specific musical lessons (see **1.3.1.5 Instruments, A musical instrument's interface works both ways**). **Rhythm tiles** form the space/time analogy: a tile's size indicates its duration; to make a piece, you have to use both note and rest tiles. **Sustain** is a player-controlled multi-pitch drone, probably controlled by a pedalboard, to give a musician (singer, violinist, etc.) a stable pitch reference against which to refine tuning.

Expressive Keyboard Harmonica. The reed organ, the accordion, and the melodica are all keyboard-izations of the harmonica; none retains the harmonica's ability to bend notes; a keyboard instrument with the expressive potential of the harmonica would be nearly on par with the saxophone as a solo instrument for blues and jazz.

Two-phase articulation. Bowed strings have a built-in two-phase articulation: up-bow and down-bow. Breathing is also two-phase, but an inhalation is seldom used to articulate musical notes (the harmonica being the notable exception). Other musical instruments could be made in which both parts of the breath cycle were used for making notes, and the characteristic nature of breathe-in/breathe-out could be incorporated into musical compositions. (Other metaphors for this include go/return, do/undo.)

MIDI pedal board. To my knowledge, nobody makes a full-size (American Guild of Organists/AGO) organ pedalboard as a stand-alone unit; the best one can do is scavenge a pedalboard from an electric organ and retrofit it with MIDI.

More axes of control for keyboards. Conventional keyboards sense timing and velocity (and, with the clavichord, aftertouch). Some MIDI keyboards have aftertouch, and some also sense lateral movement (in/out, left/right). None that I know of know where they're touched, or by which finger.

Guitar fingering tool. In this program, you'd play notes on a MIDI keyboard and the software would generate a display which showed all the potential fingerings for the current set of notes; between each group of fingerings and the next would be a connecting line showing the difficulty of the transition. Using this, a non-guitar-playing composer could find chords and transitions between chords that were relatively easy, and write guitar music that was idiomatic and not more trouble to play than it was worth. Somewhat related to this is the **Violin Keyboard** (see web page) which helps composers get familiar with violin fingering.

Ad hoc musical instruments. The performer should be able to modify the instrument. This is difficult for physical instruments, but for virtual instruments, it's straightforward. A reconfigurable pitch/tonality grid would be simple; instrumentation likewise; methods of articulation are a harder problem, but not insurmountable.

2.1.3 Tuning, temperament, tonality, consonance, harmony

While tuning, temperament, tonality, consonance and harmony are all related to each other. There are several ways to approach them. The physics of resonating bodies shows us why the sounds emanating from musical instruments often have frequency components related by whole-number ratios. There's the human hearing mechanism, which makes use of these whole-number relationships to make sense of sounds in our environment. And there's existing musical compositions and performance practices, and how they make use of whole-number relationships. To fully understand how a composer combines notes with each other, we need to be familiar with all these areas.

Physics and Psychoacoustics. I still haven't done a visualization of the physical phenomena that give rise to whole-number-multiple harmonics; I ought to. Once you know what causes harmonics, it's easy to calculate where they are, and I've done some displays to let you look at them. **V17** shows the harmonics in a linear pitch space and **V21** combines a linear pitch space with a linear frequency space. **V22** is like V17 but has a pitch space related to the ear's **critical band**. Experiments in how to display timbral energy in critical-band-related ways are continued in **V23** (splashy), **V24** (more subdued) and **V26** (adds information about beating); next to do here is to make a version that shows where the various pitches fall. **Lisajeux** is a standard Lissajous-figure generator (note: in the interest of full disclosure, I must point out that its seductive display is misleading if you want to understand the primary origins of consonance and dissonance). **INTERVAL** is a slide-rule-like tool for examining how the harmonics of two notes interact.

Intervals. The interaction of whole-number-ratio harmonics give rise to different classes of intervals between notes, each class having its own characteristic sound. **DYAD** shows the interval types between notes; variants include: **CONTEXT3** (shows notes in multiple octaves in a grid), **CONTEXT4** (shows notes in octaves in a spiral), **CONTEXT2** (shows pitch vertically, time horizontally, interval types by connection color; a.k.a. "scrolling ball of yarn), **SANDYAD** (tuned for VCR recording), **V16** (slightly different version in the V-series environment. Western harmony focuses on the intervals of the 3rd (major or minor) and 5th; **V19** shows a hexagonal lattice (like Mathieu's example 22.1, page 64).

Tonality/Circle of Fifths. Around 1990 I started working on how to visualize relationships involving the circle of fifths, and discovered **Harmonic Coloring**. This was incorporated into the **MAM**, and presented in a circular display (**WHEEL**, and later **V18**, but it still could stand some improvement; see also **ColorWheelNarration.doc**). It's sometimes useful to consider other circles, such as the **Circle of Pitch Families**. In **V12** I expressed the color wheel as a staff (circle of fifths are the 12 staff lines, horizontal position is absolute pitch); the motion of the tonal center in this display led to **V13, V14** and **V15**, which combine the circular display of "harmonic direction" with the idea of "harmonic weight." I should do a video example of this showing a quick, and then a slow, modulation.

Tuning/Temperament. Once you decide how many distinct pitch classes there are (in western music, 12 is a common answer), you're faced with the problem of tuning, since you can't tune all (12 * 11 / 2 = 66) possible intervals to whole-number ratios. **TANTRUM** (and its nice MacOS version, **Tantrum**) provides a way to look at 12-pitch tunings (this is a natural for a Java version), and **JUSTUNE** extends this (though not as conveniently) to pitch sets with more than 12 members. **DIVISION** is useful if you want to see why we settled on 12 pitch classes (rather than, say, 11 or 13), and shows the tradeoffs in other numbers. The chart **The Family of Equal-interval Tunings** shows the interesting relationship between Pythagorean, Equal-tempered, and Just tunings. I've admired the **Clayson and Garret Tuning Compass** and have considered stealing aspects of its design to make some kind of tuning aid.

Incomplete. Things related to harmony that I've wanted to work on but haven't done anything remotely successful (or, in some cases, anything at all) with include: **Visualizing the function of harmony** (for example to specifically visualize the pull from dominant to tonic, the effect of leading tones, etc.), **visualizing Schenkerian scores** (if you could display harmonic functionality, you could then have tools to summarize a score to various degrees), **display of bitonality** (though to a certain extent you can see this in the MAM display with Harmonic Coloring), and **showing dissonant "overlap"** (in **MATRIX** I tried using parallelogram overlap for this, but it wasn't too successful).

2.1.4 Rhythm, time

Rhythm is perhaps the most important part of music, which makes my failure to come to grips with it all the more signal. In **Furry** I tried doing an FFT-like (Fast Fourier Transform) transform on rhythm, which gave a little insight into the question "how fast is this music moving" but it was otherwise unsuccessful. Since rhythm involves repetition, I considered the possibility of showing "returning to the same place" in a repeated rhythm as related to "returning to the same point on a circle"; in **CONTEXT7** and **V08** I experimented with this, and in **CONTEXT8** I applied the idea to right-angle changes of direction (though the result seemed more related to **Broadway Boogie-Woogie** than to what I'd hope to show).

Sometimes, the rhythmic structure of a piece lends itself to a specific visualization. James Tenney's **"Wake"** provided a wonderful vehicle for a visualization (I did several linear versions and several radial ones, leading eventually to **V01** and **V02**). I should make some more copies of this to give to interested people.

In **V04** I tried to make a display which showed deviations from a completely regular rhythm, as an aid to learning to play scales smoothly. It would be nice to develop a display to show the differences between idealized note lengths and actual performed lengths.

Will Pitkin's project (**visual phonology for deaf infants**) reminded me to think about the non-linear nature of psychoacoustics: the "now" moment is not like other moments; in **V25** I tried a non-linear time display to emphasize the now moment.

My best lead for how to approach rhythm came from reflecting on how satisfying it was to use the Conductor Program; my rhythmic "walking" through the piece captured a surprising amount of the feeling of the piece. So I'm contemplating building a **Rhythm "animal"** which "walks" through a piece, with various limbs corresponding to the various beats (and smaller sub-limbs to handle the sub-beats; see **Gait.doc)**.

2.1.5 Sound analysis, DSP (digital signal processing)

Several of my projects require some kind of sound analysis, so I need to develop a DSP software toolkit for getting data to/from (typically Windows/SoundBlaster or MacOS) sound hardware, extract FFTs, etc. Originally, I saw my sampler as playing a big part in this, but the DSP abilities of a PC are probably adequate. Waveform and FFT data could be used to implement the **Voice Tracker**, or to add envelope data to MAM (or other) displays (or to develop a library of envelopes to use when the MAM data comes from MIDI). Timbral data could be used to make a **tool for piano technicians** to help them get a piano's tone more uniform (or make one piano sound more like another).

2.1.6 Pedagogical tools

Rhythm cards is my best pedagogical tool to date; in it, the student's gait becomes the metronome against which rhythms are performed and evaluated (see **rthmcrds.htm**); I should prepare article for Journal of Research in Music Education on this.

Byzintune is a tool for students of Byzantine chant (see **byzintun.html**). It includes features to help with the notation, the scales, etc.

There have been various programs called **Music Construction Set** (including one for the Sony PlayStation); I'd like to do one myself.

There are various **Music Education Toys** I'd like to work on, especially ones that have the engaging quality of video games. One that's fairly close to the front of the queue is a **sight-reading game** that would show you notation and you'd play along; it would make up music that was just at the edge of what you could handle (which it would determine based on how well you were playing along). An early version of this would be a toy which required that you do something regularly (once a measure) while doing other tasks in between.

2.1.7 Fluid motion

Perhaps the most inspiring music visualization for me is **Oskar Fischinger**'s early black-and-white hand-drawn charcoal animations (see videodisc). I started playing with these in my **Oskarettes**, but there is still a tremendous amount to explore (next up is a general-purpose arrow generator). My earliest music visualization experience involved **dancing notes**; I would like to recreate that. A visit to the Monterrey aquarium reminded me that **Jellyfish** would serve well as music visualization dancers: they can be shaped like just about anything, are free to move in almost any way (no bones to limit their motion, and could have muscles everywhere), are colorful, have lots of parts (to do multiple things at once), etc.

In addition to abstract motion related to music, there is also the flowing motion of the performer. In **V06** and **V07** I experimented with using a touch-tablet as a way to get "bowing" information in a computer; this was combined with pitch information from a MIDI keyboard to create a virtual 2D bowed instrument. However, to make this really sing, I need to add timbre modeling related to the bowing direction, speed, pressure, and keyboard strike velocity.

2.1.8 Timbre

My one experiment in timbre is the **Voice Tracker**, a tool for visualizing the singing voice. I should do a better presentation (compare/contrasted it with other voice displays, spectrograms, etc.) and make a real-time version.

Musical **texture is related to timbre**, in that both are involved with the distribution of spectral energy (timbre in a small time window, texture in a larger window). To understand musical texture, it would be good to have tools which let you examine complex timbral events over time.

2.1.9 Specific pieces to visualize

Most of my visualization techniques are designed for general use, and can be applied to a wide variety of pieces (though, depending on what's going on in the music, one form of visualization may be more interesting than another). But there are some pieces that I'd like to do custom visualizations for. In **Nagoya Marimbas**, two instruments follow similar paths through evolving sets of notes. I'd like to do a visualization in which the sets were mostly static (so that it was clear how they were changing—perhaps having the changing notes move from one position to another) and the two instruments were "creatures" moving through the sets. **Beethoven's Bagatelle opus 33 #5** could be visualized as a pinball machine being played. I've started thinking about a visualization for Bach's **Brandenburg III, first movement** in which each of the sections (on the order of one to ten measures long) has its own "mechanism," and there were smooth transitions from one to another. I'd also like to do a John Adams piece, and perhaps some by Richard Grayson.

2.1.10 Music visualization environment/toolkit

My first MAM implementations (1985) was a single-purpose program: it just played a single piece. In the second version (1988) I made it somewhat more flexible, in that there was an application environment in which various operations and configurations could be selected. In MAM4 (1999?) I went a little further toward having a general-purpose visualization toolkit, in that multiple applications could coexist and share the same MIDI data. It would be nice to push a future MAM further in this direction (see **MusicVizToolkit.doc**), to make a flexible visualization synthesizer. If the Sony PlayStation supported MIDI I/O, it'd be a perfect platform; even as it is, it would be a good environment (MIDI could be brought in via MIDI files). As a way of deciding what "primitives" (basic graphics operations) to include, I'd like to mine the visual languages of Fischinger, Klee, Kandinsky, Cuba, Mondrian, etc. and perhaps develop an **animation description language** in which it as easy to specify effects in these idioms.

2.1.11 Other music-related projects

Related to the animation description language is a **performance description language** in which one would specify details of a performance. Manfred Clynes has done some work in this direction; his system specifies small details in a fairly global way; I'd want to develop a system by observing what teachers in a master class tell their students (and figuring out what their suggestions meant).

It would be fun to use **music visualization in conjunction with mechanical musical instruments.**

The **SQMIDI** program generates squares in response to MIDI; this is simple, but it is still fairly expressive, and could be improved lots of ways. **V10** and **V11** both have falling rhombi; this is somewhat effective, but more could be done to improve the "history" of the display.

The **VOICEIZE** program assigns MIDI note events in a real-time incoming data stream to MIDI channels, based on pitch level, contiguity, etc. allowing a performer to play contrapuntal music on multiple instruments from a single keyboard. This program could be made much smarter.

People have done some work to **translate gestures** (e.g. from a mouse input device) **into notes**; there is still a lot that could be done in this area.

There are already some **visualization modules for Winamp and iTunes**; it would be nice to make some of these myself (see **iTunesDeal.doc**).

Adding text to visualizations would be useful in lots of ways: titles, lyrics, music theory, fingerings.

My first **click-track tool** was pretty primitive: it would just do some number of beats at a given tempo, then another number of beats at another tempo, etc. A more full-featured tool is easy to imagine.

The **CIRCLE** program is a poor-man's way to make the Sheperd tones illusion. It'd be nice to make a soundcard-based version of this.

A **musical instrument is a collection of possibilities**; it would be good to develop a picture which showed the range of possibilities (in addition to what was happening at a specific moment).

There are **music recognition systems** being used to automate on-air song recognition (see **Toureville** article); it would be interesting to study these, to see what information they extract, and how they organize and compare it.

In pre-computer days, the primary function of music notation was for a composer (or arranger, etc.) to communicate with a performer. As a result, it was important for music notation to be standardized. However, a computer could easily translate from one notation system to another, so it's time that we developed **configurable notation systems**, which a student could configure to his/her preferences.

2.2 Non-Music-Related Inventions, Software, Other R&D

2.2.1 Language, speech, writing

Speech visualization. The best tool we have currently for visualizing speech sounds is a spectrogram. To improve on it, we'd need to consider our mental model of speech, and make a display which maps to simply, but coherently (see **Pickover**'s thoughts on representing speech syllables). A goal here would be to make a display in which one the display of two people's voices looked as different (and characteristic) as they sounded. An improved tool could be used by the **deaf**, for studying (foreign) language, etc. I think of such a tool as a **phoneme animal**, an "actor" that does what the mouth does, but turned inside-out, so that instead of shaping hidden resonant cavities, it would show you what resonances it was shaping, what pitch level it was creating, what fricatives it was adding, etc. There have been software tools to **visualize the speaking mouth in cross-section**, but these are fairly hard to look at. Another approach is to describe/visualize words in terms of closeness (in some global "word space").

Conversely, I'd like to make a **one-handed phoneme generator** to create speech sounds. Is the bandwidth of control we have of a single hand enough to carry the information in an expressive, inflected speech stream? My sense is that it's close, though my first designs don't come close to being adequate.

Written language is largely based on spoken language, but we've developed other forms of communication that are based on sight rather than sound (whether they could be called "languages" is a matter of opinion): drawings, paintings, movies, etc. The problem with these is that it takes prohibitively long to draw a picture of something. Sign language is faster (and has some features of movies), but is constrained by an essentially linguistic structure. If a computer mediated our picture making, could be create simple movies (cartoons?) in real time? Could we develop a completely **graphical language** (see **GraphLang**)? Just in case, I've bought **graphlang.com**. And in general, the **future of language** is worth contemplating; after the industrial revolution, people still walk (and ride animals that walk), but they do other things that replace walking; will the information revolution provide things which are as revolutionary as the jet airplane?

Tense in language is a way of encoding the relationships among the speaker, various past, present, and future events, conditions, desires, intentions, and possibilities. These relationships could also be represented graphically (with a time line and symbols); such a representation could be used as part of a graphical language, or as a way to study the differences in tense between languages (by providing a lingua franca that was consistent, comprehensive, and not biased toward any given language).

Could the **structure of an argument** be represented schematically? (See **Making Discourse Manifest** web page and **Ask A Spec** in the **New Ideas** folder.)

The standard approach to calligraphy is to have "letter forms," which are then embodied in strokes in various ways. My idea of how to do **computer calligraphy** works differently: letters are not forms, but methods—which is to say, programs. Attributes from the environment (nearby letters, mood of the writer, quality of the paper, etc.) would affect how these programs operated; and they would operate somewhat concurrently (though to the greatest extent during the laying down of the line directly involved in itself).

The program **V05** creates a **parade of letters** that a performer can lead around the screen. I'm not sure what more could be done with this.

Wiggle Word is my name for software (not yet written) which does morphing between letter forms. The idea is that it would be used to make poetry in which words changing into other words played a part.

Folly is my video performance art setting of a Samuel Beckett poem, expressed as a continuous mental editing session.

2.2.2 Computer, Internet

Net Release (see web page) is my idea for solving the problem of royalties in cyberspace. It's too radical to be adopted whole hog; could it be modified to make the transition to it more palatable to people with a vested interest in the current paradigm?

Thing like my **RevLink** idea (reverse hypertext links) have appeared, but they haven't caught on. Microsoft tried adding something like RevLink to their browser, but since it was Microsoft-centric in its benefits, it was rejected.

2.2.3 Driving

When I was contemplating living in Los Angeles, my plan was to write a **book on driving** (see Pendaflex files on **Driving** and **Traffic** for articles). It never happened, but I'm still enthusiastic about it; subjects would include: safety, traffic flow, the statistics of accidents (and near-accidents), automobile legal issues, sociology, tips and techniques, road and highway design, driving stories, and the psychology of driving, car choices, etc.

Car designers are working on automatic collision avoidance for cars, but it's hard to design an automatic system that's reliable enough that you'd want it to take over for you automatically; my preference would be to design an intelligent **automobile proximity indicator**. Sensors would detect the distance between your car and the car in front and in back; the information from these could be presented visually (red=ahead car is closer, green=behind car is closer, brightness indicates how close, flashing indicates changes) or aurally (less than octave = ahead car is closer; greater than octave = behind car is closer; loudness indicates closeness; pulsing indicates changes).

When I first starting thinking about **visualizing traffic flow**, few people were working on it. Since then, articles have appeared in lots of scientific and mathematical journals. Most of these are for studying traffic as a dynamic system; less common are systems for displaying flow/congestion information to commuters and travelers.

Automobile expression. Part of the reason driving is as anxiety producing as it is is that drivers are not able to communicate with each other. What if drivers could talk to each other (either directly, or by having their speech turned into text and presented on an electronic billboard), or exhibit various kinds of body language through their cars?

2.2.4 User interface

We've only scratched the surface in the design of **user input devices**. See Pendaflex file **Interface** for lots of interesting experiments. Something that's under-explored is the design of the user interface (hardware and software both) to allow users to examine complex data sets. We need to develop a **language of manipulation and examination**; we have some basic things like Select and Drag and Zoom In/Out, but there are no standard ways of doing Sort, Compare, Find Correlation, Extrapolate, Simplify, Filter, etc. Part of the problem is that our way of "touching" a computer — with a mouse — is so limited. Consider what it's like to play with a piece of clay, and then try to imagine doing that on a computer with a mouse. We need more contact, either by expanding the bandwidth of the input/output devices, or by supplementing them (e.g. with speech commands). Music, being so multidimensional, is especially in need of improvements.

Related to the language of manipulation is a **gestural language of editing** that could be part of a **high-performance editing interface**. When Lisa started doing medical transcription, I thought about a GUI (graphical user interface) for doing transcription quickly: you'd hear the text spoken and see the computer's attempt to translate it to text, presented on a graphics tablet (like the Wacom); you'd work with a stylus in one hand and the other hand on a command keyboard; editing would consist primarily of two types of action: touching an incorrect word (which results in a popup containing likely alternates), or selecting a formatting command (indent, paragraph, underline, etc.); only rarely would you type on a QWERTY keyboard. (See **Glove Edit** in **New Ideas** folder.)

Typing Master Class (see web page) is my idea for a "graduate level" typing course; most courses take you from 0 to 60 wpm; TMC would take you from 60 to 120.

My **multi-resolution slider** design solves the problem of specifying and adjusting non-integral quantities over a wide range of values (0 to a zillion) with speed and precision. The slider display looks like a car odometer; when you do a mouse click on it, it becomes active; as you drag the mouse up/down, the value changes smoothly; the amount it changes varies based on the left/right position of the mouse. This kind of control is especially necessary when you're experimenting with musical pitch, which varies over a wide range (many octaves) but has to be specified with great precision (on the order of 1/1000ths of an octave).

At some point I concluded that a user interface that presented the user **knob plus two keys** offered a nice balance of simplicity and power: the knob gives you continuous control, and the keys give you articulation. The work here is to design an interface that's comfortable and very responsive, and to see what circumstances it's suited to.

Now that DSP technology is cheap, my **MIN/MAX volume control for noisy environments** is feasible. In a normal volume control, you set the gain. In the MIN/MAX control, you set the compression curve. Instead of a single knob or slider (-----|-----), you have two (----|----|----); one controls the maximum the volume will ever reach (compression kicks in near that), and one specifies the noise floor (expansion kicks in near that). The typical use would be in a car, where you don't want the music to drown out conversation, but you'd want it to stay above the noise level when there was anything worth hearing. An enhancement would adjust the noise floor based on ambient noise. I should either patent this or put it up on my web site.

The **double mouse for 3/4D control** has been invented by others, but I haven't seen one in operation yet.

The **Colby** one-handed typing interface is cute; I should write some freeware **training software** for it so that people who were interested could practice it.

2.2.5 Other

Round-robin pecking order. In many situations, we need a way to equitably resolve questions of priority. For example, when two cars arrive at an intersection, there's the question of which one goes first. The normal "pecking order" is deterministic, but unfair to those with low ranking. If the ranking is made circular and priority positions are assigned so that no two points 180 degrees apart from each other are assigned, an equitable ranking is established: every driver will get to go first approximately half the time.

Hearing AId is an AI-assisted hearing aid. There are various things that could be done. Phase information could be used to determine what sound source the subject was interested in (which way the head was facing), separate out the phoneme stream from that direction, then resynthesize the signal to emphasize that stream (and possibly enhance it to make up for frequency-specific hearing losses), cancel signals from other directions (when they conflict), and shift the effective position of signals from other directions (when they don't conflict). If the subject's hearing loss in some frequencies is too great, the unrecoverable phoneme information could be shifted into another range (again, with canceling to reduce aural clutter).

Water Wall is a fluidic-oscillator-based shower that provides a moving stream of spray without moving parts (other than the water; see notes in **New Ideas** folder).

There are popular **diagram**s that show a high-level the **flow of money** in the world; this could be augmented by adding the flow of **information, goods, services**, and other things **between consumers, producers, advertisers, media**.

A network of nodes can often be conveniently flattened into a 2D map. However, when two such networks interact (e.g. in telecommunication systems, where each network represents the state machine of a communication channel endpoint), you need to think about two points, one on each network. One way of doing this is with a **ball rolling on a map** visualization (see **Ball on Map** in **New Ideas** folder).

To discourage graffiti, paint buildings with a **fractal pattern** that minimizes the visibility of new markings regardless of their scale (see **Fractal Coloring** in **New Ideas** folder).

Human beings have a very limited memory for certain kinds of FILO (first in/last out) information; a **game which builds your "mental stack,"** if practiced from an early age, could possibly help us push our natural limit (see **"Stack Game"** in **New Ideas** folder).

Boxes and lines (a.k.a. Bubbles) is a generic graphical environment with real-world properties. In the simplest implementation, this would be used for doing flow charts (which would auto-organize so that you spent time on content rather than pulling things around) and other organizational aids. More sophisticated versions could help with building design, sentence ordering, web design, etc.

The most tantalizing application of **3D density sensing** has got to be: to feel inside a patient's body (see **3D Density** in **New Ideas** folder). A single point approach (a 3D mouse) would be fairly straightforward to implement, but better would be full-hand feedback (possibly through microscopic programmable density-altering foam/beads).

It might be worth exploring exactly how it is that **people can tell a 3D image isn't real** when it isn't.

There are various tools for organizing events in time (calendars, Gantt charts, 1D and 2D timelines), but a general-purpose **History tool** has yet to be developed. You'd want it to have the basic time-line function, but be able to alter your view by context, filtering, subject, overlaps of various peoples lives, examination/extrapolation of a single life, etc.

I made a crude tool for **visualizing stopwatch rhythms**, but it didn't have any features that the ones that jewelers use didn't have (and which a computer implementation could easily support), to help you see various common anomalies.

Animate Jaacov Agam's sculpture at the Honolulu Hilton. This is a wonderful, huge, metallic sculpture. I wrote the Hilton about this project, but from the answer it was clear that I'd need an in with somebody higher up.

The **camera lucida** was widely used by artists after perspective was discovered, but has not been used much since then. I have a lovely one (made by Holbein), but it's clear that it's in need of a **high-tech redesign**.

I'm fascinated with the idea of **MIDI control of lights**. I don't know if there's anything interesting to be done, but it would be fun to experiment with this.

2.3 Presentations (of things other than those listed under specific projects)

2.3.1 MAM-related

MAM web site. The MAM home page is in serious need of simplification and reordering (see **mam.html.doc** and **mamSiteLayout.doc)**. The MAM time-line should include a lot more items, and have thumbnails for each subject (see **MamTimeLine.doc**), and there should be an exhaustive list of developments. The MAM video page should have thumbnails of all the pieces on the tapes. And there are several items to add to the **MAM FAQ,** such as: Could the MAM be a tool for composers? (see **FutureMAMFAQitems.doc**) How are MAM renditions created? (see **MAMFAQtemp.doc**) What should I look for in the MAM animations?

MAM museum exhibit. The idea is to have a touch-screen that controls the selection of pieces, speed of playback, volume, scrolling, and other parameters (see **Museum Animal**).

MAM in film/video/animation festivals. I should do some pieces with an eye toward this (and even if not, get the gear assembled so that I could if I wanted).

2.3.2 Other: music-related

There's a lot to say about **Visualizing Tonality, Harmony and Interval Relationships** (see **CircleOfFifths.html.doc** and **"Experiments in visualizing musical harmony"**, **ColorPitchAssociations.doc** and **"The Problems of Synesthesia"**).

I should **transfer music visualization examples to DV (digital video)** for presentation on-line (and perhaps DVD, see **MusicVizDV.doc**); this includes the CONTEXT-series and the V-series, plus all the others.

The frequency of light is a poor analog to the pitch of sound because we sense light frequencies by interpolation (and hence can't separate frequencies out), but sound frequencies directly (and thus can separate them); this doesn't mean that color can't be used for other things, though, such as pitch <u>class</u>; see **LightSpectrumInappropriate.doc**).

The Materials of Music is my working title for a video/web presentation that explains the basics of musical acoustics, psychoacoustics, and music theory from the ground up (see **TheMaterialsOfMusic.doc**).

French Horn Stopping (see web page **hornstop.htm**) is my explanation of what's happening when French horns change pitch due to stopping the bell with the hand or mute.

Contrapuntal keyboard exercises (see web page **grabbag.html**) is a collection of exercises to help a student play contrapuntal music.

The Failure of the Theremin is an (as yet unwritten) article on why the Theremin, beautiful instrument that it is, was such a failure (due to the lack of force feedback, etc.).

A good demonstration of harmonic motion would **use a simple folk song to show I/IV/V motion.**

The **Cutting Edge Razor Band** page should have a selection of the band playing.

I should write a theoretical overview of **harmony**, touching on interrelation of psychoacoustics, function, tuning, psychology, progressions, timbre.

I should **think of something useful/fun to say at Tufte seminar** so that if I attend again I don't make a fool of myself. Likewise, **YLEM** is a nice forum (which focuses on the intersection of art and technology), and I should prepare another presentation for that.

The interaction of tuning systems and the kinds of chord progressions that often happen in music explains **why choirs go flat**; this would make a nice web page (with an animated picture showing the nub of the problem).

There are lots of exercises one can do to improve one's **sight-reading** skills (see **SightReading.doc**); I should collect the ones of these I've invented so far and make them into a web page.

My **non-MAM-related musical inventions time-line** should go on line (see **MamTimeLine.doc**), and possibly be made into a video.

2.3.3 Other: non-music-related

My **non-music-related inventions time-line** should go on line (see **New Ideas** folder), especially the **world's best back-scratcher**.

Many people don't understand the difference between pitch and resonance (the reason you can tell that a child is speaking and not a woman has to do with resonance, not pitch); to **show formant shift, sing "EEEEEE" while helium is increasing/reducing the amount of helium** (or simulate this electronically).

With the right animation, it should be a cinch to explain **why the color wheel is a wheel** (show the 3D color space, cut slices, etc.).

I've done some funny things in the area of **genealogical charting** that might be nice to share.

Lloyd Watts has done some nice diagrams showing the **structure of signal pathways in the ear**; I think they could be improved with more detail (and clarifying the dataflow).

I wrote an article about **SIGNature**s, which might be nice to put online.

Computer jokes. When I encountered the first standardized computing environment (DOS PCs), I wrote various practical joke computer programs; I should document these, and bring them up to date (e.g. the "Belly" button in a GUI).

Any logician knows that **(A and B) == not(not A or not B)**, but most non-logicians don't. With the right diagram, it's easy to see that they are identical, that they're just two ways of talking about the same picture.

Surprisingly many people believe **the myth of 3D**, namely, that 3D gives you "a whole nother dimension to work with" (in the same what that the transition from 1D to 2D does); this is not true, and 3D visualization is a waste of resources in many contexts (see **McCutchen** letter for the beginnings of this argument).

Why do mirrors reverse left-to-right but not top-to-bottom? There's an answer to this, and my **reverse visual aid** helps people see it.

2.4 Promotion, Marketing

Perhaps the most basic question about promoting my visualization projects is: **how do you market something nobody's ever seen?** Related to this is: **what is the MAM good for?** (A question which may or may not have a good answer.) While a certain percentage of people are thrilled to see the MAM, few people have opted to buy a videotape sight unseen (the signal exception being that Edward Tufte's customers have bought them based on his recommendation). There are several promotional avenues suggested by this: **TV promotion** (currently some of this is happening via Classic Arts Showcase—and could be expanded, but there are other ways, including Jack Ahearn's approach), **web promotion** (so far there's only a single, poor-quality QuickTime movie), **follow-up sales** to customers who've seen it before (easy, since I have a mailing list of Tufte's customers). There are various **film rating services** which could add the MAM videos to their publications, and **music/education catalogs** that might want to carry them (though my experience with these in the past has not been good). Another possibility is a **freeware MAM** program; this could be limited in function so as not to compete with videos. If it were done as a slick shareware version for MacOSX, it might be easy to get it demoed in **Apple's new stores**.

There are some kinds of promotion which do not generate sales directly but which get the MAM into places where it would be appreciated. The current offer to **give free videos to educational institutions** has gone a small step in this direction; this could be beefed up by **actively promoting to music history/theory/appreciation teachers** (and I should work with Carlo Caballero to develop something interesting)**,** to **schools**, and to **music education societies**. The **tuning software** and some literature could be sent to people who are active in the area of tuning. There are many people who've expressed an interest in the MAM; it would be nice to send them complimentary copies of the History video, descriptive and promotional literature (see **MAMdescriptiveText.doc**), etc. I should also send a copy of the History video to **Clifford Pickover**, since he's shown an interest and shares his interests with a lot of people. And in general, it would be good to assemble **email and snailmail mailing lists** so that it's easier to send things out to people.

To market the video through **commercial video outlets**, I'd want to **improve the packaging art** and **put out versions on DVD**.

Many people have suggested that I **apply for NEA grant** (though it's not clear what for).

2.5 Things to study

There are various things it would be good for me to know more about in order to do a better job presenting my work. The colors in my videos have long been problematic; it would be good to learn **how color signals work in recording**, so as to get the best possible result.

Since color figures prominently in what I do, I should learn more about **color vision** (both in terms of science of color perception and analogies between our color sense and other sense).

3. Resources

3.1 Charts, Lists, Descriptive Material, History

3.1.1 MAM-related

Lists. There are various lists of MAM-related things: **places the MAM has appeared** (see **MAMappearances.doc**), **off-site demos** (see **OFFSITE.doc**), various **milestones (see MAMmilestones.doc**), **quotes and other reactions** (see **MAMreactions.doc** and **folder with student comments**), **pieces currently playable on the MAM** (see **MAMlistOfPieces.doc**), **people who want to know when things about the MAM have happened** (see **MAMfans.doc**), **reviews and awards** (see website for reviews, files for awards), **people I've been involved with in MAM-related projects** (see **MAMpeople.doc**).

Descriptive material. I prepared a longish description for **Macmillan/McGraw**; **Ted Rust** wrote a nice article (of which there are reprints), and I prepared a series of **white papers for the College Music Society**.

MAM on KRON. The MAM was on TV news in the early 90s (see **MAMonKRON.doc** and video). I might want to transfer this to digital video to distribute to people.

Unsuccessful Collaborations. Inside Music with **David McCutchen**; video with **Tra-La-La**; **Exploratorium**; **David Rothe at Chico State**.

3.1.2 Not MAM-related

When Sandy Cohen was teaching his class on music visualization, I wrote him a letter in which I discussed the **dimensions of sensations** (see **Cohen letter 1995sep28**). I've collected various articles and pictures of interesting **music interfaces** (see folder under **Interfaces**). It's good to keep in mind my **areas of interest and expertise** (music theory, music composition, music performance, physics, math, software, electronics, acoustics, visualization, audification, writing, teaching). Of course, I have a list of **people who do music visualization** (see **MusicVizVideos.doc**). And I have a very comprehensive **chart of names of colors.**

3.2 Environments, Technologies, Tools

3.2.1 Environments

It's useful to consider the situations in which the kinds of things I'm interested in have application. The **home** is a place where people can watch music visualization. The **classroom** is another obvious venue (including **music theory classes**, where toys for playing with acoustics and tuning would be good). Many of my inventions could be embodied as **toys**, either for children or for **old people**. Visualization of sound could be useful for the **deaf**, and hearing is a **health issue** for everybody (**hearing protection**). Exhibits could be set up in **art museums** and **science museums**, and presentations could be given to a variety of groups, including **YLEM**.

In the realm of performance art, there are many possible venues. **Nightclubs, music concerts** and **raves** are a natural place for music visualization, especially if they use **computer-controlled pianos** or other electronic musical instruments.

3.2.2 Technologies

As with Environments, there are many technologies that either I'm interested in in themselves or which might have application to my projects, including:

- **Hypertext** (the web and other forms)
- **DSP** (both software and hardware, its ubiquity makes it ever more relevant) including **Chirplet**
- **Java** (which allows things which are best expressed as software to be shared widely)
- **MacOSX** (important right because Apple is eager to get it off the ground, and new applications have a better chance of getting promoted)
- **Windows 2000** (because it's what the "rest of them" use)
- **Biofeedback** (actually, I'm less interested in this than when I put it on the list).

3.2.3 Tools, other resources

This is a list of things which it's good to remember are available:

- **Geometer's Sketchpad** (handy for doing animated sketches of certain kinds of things)
- **Casio FZ-10M samples** (though, with the improvements of DSP in computers, less and less relevant)
- **On-line MIDI files** (as material for MAM and other visualizations)
- **McGill organ samples** (if I want to do more Bach organ music)
- Roland makes a **cheap MIDI keyboard** (would be nice to package as part of a music visualization toolkit)
- **US Video VGA card** (after ten years, still the hardware with the best signal)
- **Wacom graphics tablet** (for lots of interactive experiments)
- **Screen Dance** (contains an inexpensive spectrogram that's sometimes useful)
- the **Sony Playstation** is inexpensive, powerful, and has decent graphics
- **PianoDisc** has better real-time response than **Disklavier** and is thus better suited to the conductor program
- **Ross Hippley's C64 program** is cool, and I should learn how to do demos with it

3.3 Other Materials, References, Notes

Things which I might want to take advantage of someday:

- Joseph Spencer has said I could use **Concerto Amabile** performances with MAM visualizations
- the copyright on **Milhaud's Saudades do Brazil** may be expiring soon
- I have various **3D images**
- **tactile alphabet letters**
- **Cuisinaire blocks**
- **searchable on-line patent archives**
- **drawings that are hard to translate into words** (see various files)
- **books related to music visualization** (see **MusVizBooks.doc**)
- **articles related to music visualization** (see **MusVizArticles.doc**)
- **companies worth keeping in mind** (see **MAMcompanies.doc**)
- **music-related patents** (see files)