

GlobeMusic: The Internet Scale of eMusic-Making

Max Mühlhäuser¹, Michael Welzl², Jan Borchers³, Rainer Gutkas²

Technical University of Darmstadt (DE)¹; University of Linz (AT)²; Stanford University (US)³
max@informatik.tu-darmstadt.de; michael/rgu@tk.uni-linz.ac.at; borchers@stanford.edu

Abstract

The Internet is not only the global data highway but also the upcoming fourth mass medium. In this role, it procreates new musical genres and new roles of musical content providers and consumers. In an attempt to contribute to and explore this exciting (r)evolution, we developed GlobeMusic, a novel system and concept for eMusic-Making. Major ingredients of GlobeMusic were developed and exposed to a large public in the last four years. GlobeMusic and its components mainly address the following issues: computer-human interaction, individual eMusic-making, appropriate computer network protocols, musical data formats, and cooperative eMusic-making. The paper sketches the four major evolution steps of GlobeMusic because each of these steps represents a key aspect and system component.

1. Introduction

We propose novel approaches to music-making in the light of the breathtaking advancements of the Internet. We describe several evolution steps of a system developed over the last four years. We consider the following statement crucial with respect to the Internet in its role as the upcoming fourth mass medium: *like any new mass medium, the Internet is going to foster new content types (genres)*. Readers who doubt here may be referred to the fact that for the case of video, e.g., Santa Monica has become a melting pot of ingenious people and companies who are about to investigate new content types and, of course, new ways of creating them [6] – many of them left Hollywood because they felt that this famous site was too much bound to the existing genres. Two basic assumptions of the protagonists of Internet-specific genres are worth noting: i) Internet genres will be highly interactive; ii) the boundaries between content producers and content consumers will become much more blurred. While we will concentrate on music instead of video below, we will adopt these very same assumptions for research about *new Internet-specific musical genres*.

New genres emerge from a process that is influenced by a highly dynamic and non-predictable gravity field with three poles: the (changing) taste, budget, and interest

of consumers; the (improving) limits of technical feasibility; and the (least predictable) genius and fantasy of content creators. Therefore, the present paper does not claim to “know” the nature of future musical genres, but rather reports on systems and concepts developed, experiments, and large-scale user experiences as elements of the melting pot from which new genres will eventually emerge.

On one hand, we distinguished creators and consumers of contents above, on the other hand we argued that their roles might become blurred - how does this match? The following aspects must be considered here:

- In the case of music, the three-step process of composition, interpretation, and listening has long been generally accepted. However, if a Jazz band unites for a jam session without additional audience, they are supposed to be (largely) composing, interpreting, and listening all in one. This shows that even without the Internet, modern music-making tends to blend formerly distinct roles.
- The Internet is not only the first mass medium to be interactive, it is also the first one which goes beyond broadcast by embracing person-to-person and group-to-group communication. This fosters the move from merely consumptive to partly creative “consumers”.
- Envisioning the Internet as an interactive medium for the masses, we face dramatic implications: couch-potatoes and other mere consumers are supposed to become interactive, even creative! This is only possible if the interaction/creation process is easy to carry out yet highly attractive (fun, motivating, thrilling, ...) – readers who doubt this argument in general should note that i) interactive computer games do represent a booming *mass* market; ii) the authors gained very positive experiences exposing average users to interactive, even creative musical Internet-Instruments. This implies that the “producer” side in the media business moves from *authoring read-only media* via *creating interactive media* towards *software development for ‘create-your-own’ toys/tools* – in classical music lingo: “the users embrace composition, interpretation, and listening all in one, while the former composers shift to the task of inventing ever new (software/hardware) instruments”.

Fig. 1.1 relates the four essential parts of GlobeMusic to a scenario of several musicians connected to the Internet (potentially from different corners of the globe) and to a distributed computer-based infrastructure (depicted with

a “hard disk” symbol) for cooperative eMusic-making. In the remainder of this paper, we will describe these four systems and approaches that we developed over the past four years. Thereby, the steps towards the final system are coined with different “proper names” since these steps were exposed to the public as ready-to-use systems for experimentation. The steps (and systems), in chronological order, were called WorldBeat → NetMusic → NetScore-RT → GlobeMusic.

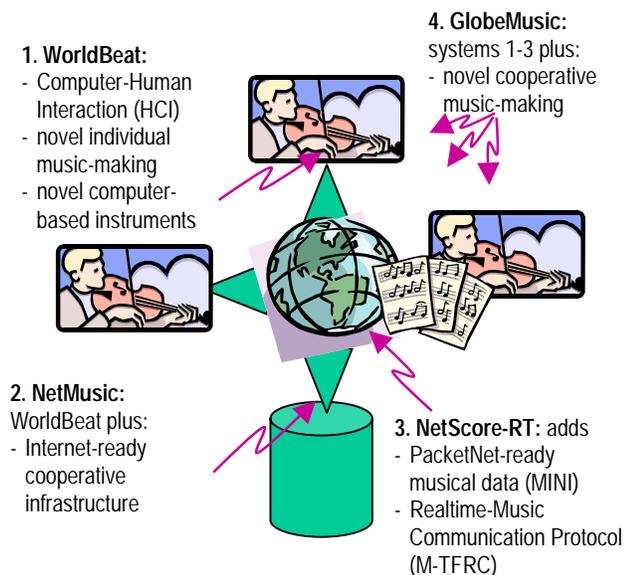


Fig. 1.1 GlobeMusic components

2. Worldbeat: removing the access barrier

Networked music creation systems aim at creating an online community of users that engages in social interaction by cooperatively playing, improvising, and performing music. Therefore, their success depends critically on a broad user base. The first step to building a GlobeMusic system then is to ensure that it appeals to a broad body of potential users across the global network.

In classical desktop computing scenarios, such as running office applications, users are more inclined to accept a steep learning curve in exchange for a long-term payoff in work efficiency. In systems that aim primarily at supporting social interaction, such as GlobeMusic, however, those steep learning curves have to be avoided to give first-time, short-time users a chance of success in their interactive experience. Ideally, such a system should scare off neither computer novices nor music novices to maximize the potential user base.

To take down the barriers that prevent potentially interested users from interacting with GlobeMusic, its user interface and functionality then have to follow a set of

rules that can be derived from the design of interactive music exhibits and similar public-access systems [2]:

Simple Impression: The user interface has to make clear that it is going to be easy to use. In particular, if it is supposed to be a system about music, it should not appear to be a “complicated application” that requires a computer expert to operate it.

Innovative, Physical Appearance: Despite simplicity, the system has to suggest (and deliver!) a novel kind of interactive experience that is more physical and “realistic” than using the standard mouse and keyboard.

Balancing Creativity and Support: On one hand, it is important for the user to receive help from the system when creating music using it. However, this guidance has to be supplied in a subtle way that ensures that the user does not feel decoupled from the creative process.

2.1 Computer-Supported improvisation

The above requirements can be fulfilled by a combination of innovative user interface technology and metaphors, with software functionality that carefully models musical concepts to assist players in the creative process.

We have developed a successful solution for the first part, supporting the user in creating input to a GlobeMusic system, by building a software component that carefully adjusts the user input in real time in order to ensure a satisfying result independent of the user’s musical “abilities”. When we exposed this part of GlobeMusic to the audience as a self-contained system, we coined it as *WorldBeat*. In general, there are three major dimensions of musical performance in which a system can try to support the user, but we will explain why we chose to control only one of them:

The rhythmical dimension (adjusting the timing of user input to match a given beat; generally called “quantizing”). While it is computationally relatively easy to “straighten” musical input to match a fixed pattern, this obviously leads to a machine-like rendition of the input. This can be improved upon by building components that take certain rhythmic irregularities into account, to create a more “human-like” quantization. Nevertheless, rhythmic support always leads to two problems: If the input is to be processed in real time, musical events can always only be delayed to a later point in time, never shifted “forward” to an earlier point. Second, and far more substantially, however, any kind of rhythmical adjustment will invariably lead to the user feeling decoupled from the creation process: If a user “plays” a note, and that note sounds not immediately but slightly later to match the beat pattern, the user will feel she is no longer playing an (augmented) instrument, but instead only controlling the process that actually creates the notes. Therefore, rhythm is not our preferred choice dimension to interfere with the user input.

The melodic dimension (adjusting the melodic patterns of input). This would mean, for example, creating whole melodies in response to a single user input event. While this technique has found wide adoption in consumer music devices (the infamous “arpeggiator” on the average home keyboard comes to mind), we believe that it again makes the user feel like he is losing a vital aspect of control over what is happening musically. The satisfaction of triggering entire musical sequences with the press of a key quickly fades, when users discover that this makes it increasingly difficult to tell the user's creative input apart from what the system is generating on its own. Therefore, we leave the melodic dimension to the user as well: He can play high or low notes, runs or chords, and the system only plays one note per user-level input event, much like a traditional instrument.

The harmonic dimension (adjusting the harmonic structure of the input). This means taking the user input, and re-mapping it slightly (changing pitches) to fit the current harmonic context. This is a very suitable dimension for augmentation, since the user can still control the rhythm and melodic profile of his performance. The system only takes the input and changes it just so that it fits into the harmonic structure of the accompaniment. This is what our ImprovisationHelper does. Another advantage of this approach is that it works perfectly even on real-time input that requires real-time output of the augmented data stream.

To understand how the “ImprovisationHelper” works in detail, we will look briefly at the software objects it consists of, and how they work together (see Fig. 2.1) [1]:

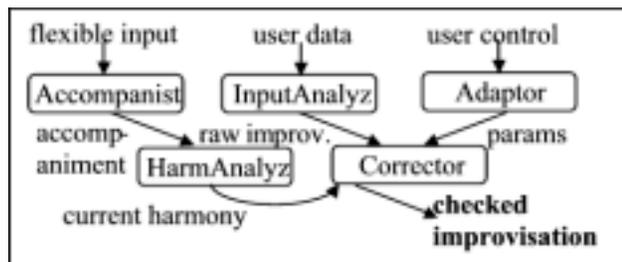


Fig. 2.1 “Improvisation Helper” structure

The *Accompanist* supplies the computer-generated accompaniment of a Blues Band (whose groove etc. can be adjusted using other patterns like the *MetricTransformer* described above). The *HarmonicAnalyzer* uses a root-parsing algorithm as described in standard music literature to determine the current chord (say, Fm^7) in real-time. The *InputAnalyzer* offers a xylophone-like playing metaphor: The user makes downbeat gestures with the two infrared batons in his hands. Gesture velocity determines volume, horizontal position determines pitch. The *Corrector* takes this input and maps it to the nearest, harmonically sound note in terms of the current accompaniment chord determined by *HarmAnalyz*.

The ImprovisationHelper has been used within a computer-assisted improvisation component of the WorldBeat music exhibit. The result is quite fascinating: People who have never before played an instrument, can walk up to the system and start improvising to a Blues band -- without playing wrong notes! This makes this WorldBeat component very attractive and popular among visitors using WorldBeat. Experienced musicians can adjust the support to either a moderately hard to play, chromatic virtual xylophone using the batons (1 octave), or they can even use a full-scale electronic piano keyboard that can be added to the system at any time (7 octaves).

2.2 Interaction and experience design

Supporting rich musical structures also requires an appropriate user interface to ensure that users can actually understand the musical concepts encapsulated in the system, by interacting with them. We have achieved that by using a set of two “infrared batons” that can be used to create musical input. They work like xylophone mallets that are used in mid-air. A tracker converts baton movement to MIDI controller data which is then interpreted and turned into musical events. The system uses the *Lightning II* baton system designed by Buchla and Associates, and adds a software layer on top of this input device that lets the user create musical input for a GlobeMusic application in a very musical way - by downward drumming gestures with the two batons (see Fig. 2.2).



Fig. 2.2: A user improvising to Blues using WorldBeat

These batons offer a number of characteristics that are advantageous to our goal of a broad user base: i) they do not look like a standard computer interface, and instead create an innovative appearance that makes the system more tempting to use; ii) they do not look like a complicated, professional-level musical interface either (such as an 80-key piano keyboard), which would scare off people who do not play such instruments; iii) they force users to

interact with the system in a much more physical way, requiring them to stand and wave their arms to play, instead of just asking the user to sit down passively and operate with icons on a small screen using a mouse and keyboard.

Since there are two batons, *local* cooperation between two players is natural and encouraged. Already in *WorldBeat*, users were often observed to share the two batons, essentially turning the human-computer interaction of a single user creating music into a human-human cooperation where the computer merely becomes a medium. This is an exciting level of cooperation in addition to the remote collaboration in *NetMusic* and *GlobalMusic*.

While it is obvious that a special-purpose device such as the infrared batons is not available to every potential user today, we believe that the trend to a larger variety of post-desktop user interface technologies is supporting our vision of a scenario where people will obtain a certain device if it is really tied in to an online experience and promises a new level of interaction.

3. NetMusic: musical cooperation

Chapters three through five are based on two simple but very far-reaching axioms: i) the margin between two sounds must usually be less than 40 ms for the human ear to grasp them as played at once (cf. [16]); ii) even considering data transfer at the speed of light and unlimited bandwidth, it is impossible to keep the delay of bi-directional transactions under 40 ms around the globe. In this chapter, we describe a basic distributed infrastructure for global cooperative exchange of musical information, interpreting “cooperative eMusic-making” as (somewhat) asynchronous editing of a common musical score; chapter four emphasizes means for reducing the real time delay between globally distributed musicians to the minimum possible, and chapter five investigates globe-spanning *synchronous* cooperative eMusic-making.

A survey [17] considered the approach for musical cooperation used by *ResRocket Surfers* the best one investigated. The underlying idea of this approach was to combine the functionality of a MIDI sequencer with the functionality of a *MUD* (Multi User Dungeon). The users could navigate through virtual music studios, in which they could participate in songwriting. Apart from the possibility to chat with the fellow participants, the users of *ResRocket* could record MIDI tracks and transmit them to fellow participants. This client-server solution can be classified as an *asynchronous* Internet approach for musical cooperation [10].

ResRocket Surfers has evolved further into a Web-compliant combination of chat room system and audio data transfer system. Using the *Rocket Power* standard, third-party systems may exchange music data over a *ResRocket* server. All-in-one, however, *ResRocket Surfers* is designed with professional users in mind: e.g., ap-

plications based on “*Rocket Power*” have complete control over all network events; and users are exposed to a large variety of features and options [10].

3.1 The NetMusic system

NetMusic extends *WorldBeat* by introducing networking capabilities; synchronous (real time) cooperation is not yet considered here. The basic data transfer scheme draws from *ResRocket Surfers*. The user interface and application design philosophy of *NetMusic*, however, followed the *WorldBeat* line-of-thoughts – in fact, *NetMusic* was integrated with *WorldBeat* and targets both experts and, in particular, novices (wrt. both music and computer use).

For the sake of simplicity, most network actions are hidden. As users click the *NetMusic* button in *WorldBeat*, an automated login and update process starts and the *NetMusic Song Page* (see fig. 3.1) is displayed. Status messages report relevant events. Novices receive hints and “how-to”s, users can listen to the (current state of the) song and go further to record a song part i.e. a track (instrument) and take (verse). In the default setup, *NetMusic* songs are configured to consist of four takes played by four instruments. The choice of instruments serves several requirements: i) predefined instruments simplify the interface for early users; ii) the instrument pattern of saxophone, piano, bass, and drums is often used in modern music and matches well with the sound of typical *NetMusic* songs, yielding a familiar pattern; iii) the instruments used sound still natural if played with a standard after-touch pattern, making them preferable candidates for the *MINI* format described in the next chapter. Each part of a song may be in one of the three possible states *free*, *in progress* and *used*. In the latter two states, songs are marked with the name of the editing site (e.g., town). This makes the users aware of the distributed nature of *NetMusic*. Unused parts host a free (recording) button.



Fig. 3.1: SongPage screenshot (“Strophe” = take)

When the user selects a free part, a lock request is sent to the NetMusic server. The server acknowledges the lock and switches to the recording screen for the desired instrument (raise conditions are resolved). When the user leaves NetMusic or if the kiosk is reset by the automatic time out mechanism of WorldBeat, the client automatically logs out from the server. A NetMusic server may be co-located with a client and is currently configured to serve up to 16 clients. Its main tasks are client coordination and broadcast, and change logging and propagation.

3.2 Asynchronous eMusic-making scheme

The asynchronous scheme described here represents a first step, complemented by a quasi synchronous scheme in GlobeMusic (cf. chapter 6). There are different ways of recording in NetMusic: for piano and bass, the user plays on a virtual keyboard using the batons. In addition to the improvisation helper function as described in the previous chapter and to the expert mode with full chromatic keyboard, we offer a virtual keyboard featuring all notes of a key such as C-major except for the fourth tone (quart). The drum recording also uses the keyboard metaphor. Several drums are mapped to a simplified virtual keyboard. The user can practice the usage of the instrument first and then start recording by clicking a corresponding button. Saxophone recording follows a special approach: the user hums into a microphone which is attached to the WorldBeat kiosk. The tone hummed is quantized to 12 semitones. In the inexperienced mode, wrong notes are mapped to the best matching ones of the key. The recorded notes are not echoed during recording because differences to the hummed note could disturb the user.

To assist the user regarding timing, a metronome is played in the background during the recording of a take. When the recording starts, the storage contents of the take are copied to a temporary storage in the recorder. This is the storage used locally while the user “edits” i.e. records his take (among others, this prevents simultaneous edits by remote users from disturbing locally). At the end of the recording sequence the take is played once again. Users can listen to their artwork and rerecord it if they are not satisfied. When satisfied with the recording, users navigate back to the song page. The server then broadcasts the new instrument/take to all clients; the clients reflect the update on their song pages.

In NetMusic configurations used for exhibits, a simple exhibition-proof scheme must be used for deciding about the “actual” song edited by globally distributed exhibit clients. The current default scheme maintains a distributed musical score either until it is complete (in the above configuration, 4x4 instrument/takes are recorded and the “last” users at each site have (been) logged out) or when a song has reached 70% completion without remote locks showing up (in this case, it is discarded when the current user logs out).

4. RtScore: quality scaling for musical scores

4.1 The network view: the M-TRTF protocol

It is a common requirement that Internet applications should adapt to the actual network conditions. This has been proven to be a prerequisite for network stability in [15]. From a distributed application programmer's point of view, the advantage of adaptation lies in a reduced packet loss ratio: it is better to use an adaptable encoding scheme and to reduce bandwidth use in congested networks than to lose a large fraction of data packets. Thus, the common connection oriented Internet transport protocol TCP has been enhanced to support rate adaptation; however, for most streaming multimedia applications, this behavior is not appropriate. UDP as the common alternative does not perform any error and congestion control at all and leaves adaptation up to the application. This openness has led to an immense body of research on quality adaptation and end-to-end congestion control as well as diverse applications with a huge range of responsiveness.

Since TCP is the prevailing Internet transport protocol, the concept of TCP-friendliness has been introduced for *non* TCP based applications. A TCP-friendly flow is defined in [3] as a flow that, in steady-state, uses no more bandwidth than a corresponding TCP flow running under comparable condition. This can be achieved, e.g., by partially emulating the behavior of TCP.

The TCP congestion avoidance policy reduces its sending rate by a factor $\beta = 0.5$ in response to a single loss event and gradually increases it by a step α when no loss is encountered. It is classified as AIMD (additive-increase-multiplicative-decrease); various other AIMD mechanisms and variations of α and β have been proposed for streaming applications [7, 9, 13, 14]. For an adaptive multimedia application, all AIMD mechanisms share the disadvantage of drastically reducing the sending rate in response to a single congestion indication.

Another definition of TCP-friendliness is that a long lived flow should satisfy the TCP response function describing the steady-state sending rate:

$$T = \frac{s}{R\sqrt{\frac{2p}{3}} + t_RTO * (3\sqrt{\frac{3p}{8}} p(1+32p^2))}$$

where T denotes the sending rate in bytes/sec, s is the packet size, R is the round-trip time, p is the steady-state loss rate and t_RTO the TCP retransmit timeout value [5]. Satisfying this equation actually means sending as much as or less than T, but sending less is not desirable because it may result in starvation when competing with TCP. Using this equation, it is possible to implement a rate-based TCP-friendly congestion control scheme that does not implement AIMD. This scheme called TFRC has in particular been shown to yield a smoother sending rate pattern [5]. For most streaming multimedia applications,

drastic bandwidth fluctuations are not desirable since they yield highly jitter-prone streams. This holds true for our application as well. Using the TFRC code from [5], we have compared simulations of TFRC and TCP across a 100 node transit-stub network (which is shown in [19] to be a good approximation of the Internet topology) with a 3:2 mixture of self-similar background traffic (following a Pareto distribution to account for large-scale web traffic [4]) and of TCP flows. It is clearly visible in fig. 4.1 that the sending rate of TFRC is indeed smoother in our simulation scenario. We therefore chose to use a custom TFRC variant in our application called M-TRTF (M for music).

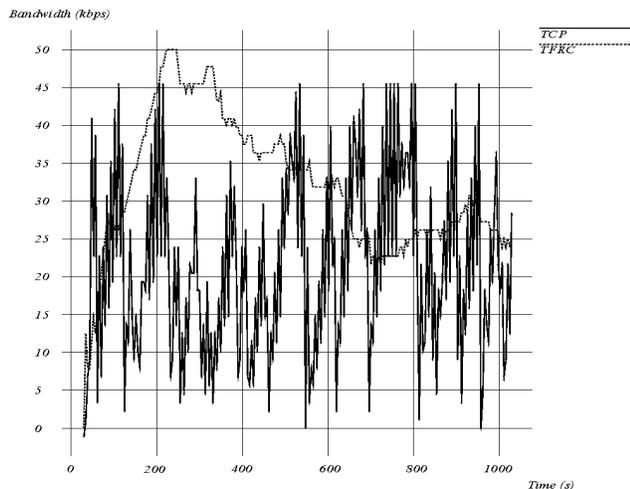


Figure 4.1: TCP vs M-TFRC sending rates.

4.2 The MINI musical format

Research on quality scaling of multimedia data has largely focused on video or audio rather than “MIDI-like” symbolic musical data where notes are represented by control data. Many published codes for symbolic music representation were created at a time when quality adaptation for multimedia communication was not a key research issue (cf. [13]); some proprietary codes may at best be suspected to include quality scaling due to a lack of publications. Work around structured audio (SA), in particular in the context of the MPEG-4 standard (see [11], [12]), has received a lot of attention in the music community lately. SA aims at combining the richness and fidelity of conventional audio codecs (such as MP3) with the compactness of MIDI. SA is design-centered around high fidelity and embraces much more than MIDI devices; as a result, its compactness will usually not go beyond that of MIDI and encoding / decoding is processing intensive. Our MINI format, in contrast, has exactly the contrary focus. Note that MINI is just a format proposal while SA represents multi-party research embracing format, codec, and more. MINI basically suggests adaptations to MIDI

that enable quality scaling for Internet. Developed before and parallel to SA, MINI should be merged with SA, in particular with efforts which address packet network transmission of MIDI like data such as [18].

MIDI was not designed to be transmitted over long distances or packet oriented networks in real-time. The assertion of a maximum transmission distance with a minimum bandwidth of 31.25kbaud (+/- 1%) and a given maximum latency did not raise any need to introduce a special symbol for chords; thus, a chord in MIDI is encoded, transmitted and possibly played back as a series of notes, fast enough to avoid audible side effects [8]. If delay is introduced between notes, e.g., due to latency, chords turn into arpeggios. Jitter makes arpeggios sound particularly unpleasant, as experiments of ours have shown.

One might question our choice of MIDI as a starting point in the first place. We took this choice because MIDI is the dominant music encoding standard and the basis of our hardware – a primary requirement was therefore easy translation from and to MIDI. Our format, MINI (“Musical Instruments Networks Interface”) is basically a subset of MIDI with special encoding of chords and notes. It can be used for what is called layered encoding: data are sorted in order of importance and given an appropriate priority. Layered encoding is a prerequisite for quality scaling i.e. the adaptation of required bandwidth (and other Quality-of-Service parameters) to network conditions, which we base on M-TFRC as described.

The derivation of MINI from MIDI can be interpreted as a three-step process as follows.

Step 1 – Saving space: a design goal being smaller size, the following MIDI messages were removed:

- *Tune Request:* This message was only designed for analog synthesizers. Even if such a device were used in a real-time long distance collaboration, the separated musicians would not notice the differences in tuning.
- *System Exclusive Messages, EOX (End of Exclusive):* These messages are for vendor-specific features only.
- *Active Sensing Message:* This message is related to MIDI errors, to be sent at least once every 300ms, hence is inappropriate for a packet oriented network.
- *System Reset, Local ON/OFF, All Notes OFF:* Similar to the Active Sensing Message, these messages are related to MIDI errors and testing. Such problems can and must be solved locally in Internet settings.
- *Omni ON/OFF, MONO ON/Poly OFF, MONO OFF/Poly ON:* MIDI supports logical splitting of communication into channels. For MINI, communication is structured according to our distributed software, hence these commands can be omitted. Note that the distinction of monophonic and polyphonic music is implicitly encoded in our chord representation described below.
- *Song Select, Song-Position-Pointer, Start, Stop, Continue:* These messages are used to control sequencers. They are not relevant in our context.

Step 2 – Supporting Jitter-prone networks: In MIDI, one must explicitly disable each tone played with a *Note OFF* message. This approach is not only a waste of data for tones which have a fixed length (e.g. a vibraphone, a piano or percussion instruments in the “regular” way of playing - special playing techniques are usually encoded as individual instruments), it is in particular infeasible for jitter-prone networks. In MINI, we use a single bit for identification: if the bit is cleared, the instrument is to be played with a fixed length. This predefined note-length period may be used to define a quasi metronome based playing scheme and may be combined with multiples (e.g., describing quarter, triplet, and half notes if the fixed length defines eighths) and with pauses – the receiving process may then buffer notes and play smoothly and accurately even in presence of considerable jitter. For compatibility reasons, we also support the note-on/off concept (in which case the above-mentioned bit is set for note-on and cleared for note-off), to be used of course over switched networks or in case of very low jitter only.

Step 3 – Layered chord encoding: There is no special symbol for chords in MIDI. Apart from the aforementioned arpeggio effect, this is clearly a waste of data. We assume that the MIDI range of ten octaves will usually not be entirely used for streaming music. We have seven bits left in the first byte of a MINI command (one is reserved for Note On/Off control, see above). To determine what we can encode with these seven bits, we recall the formula for the collection M of all sets (here: chords) with at most k elements (here: tones) as a kth order combination of n elements with repetition and without ordering:

$$M = \binom{n+k-1}{k} = \frac{(n+k-1)!}{k!(n-1)!}$$

Using a seven bit word, it is possible to describe all combinations of n and k where M does not exceed 128. This yields the following possibilities:

| Encoding | Range (n) | Voices (k) | Combinations (M) |
|----------|-----------|------------|------------------|
| 1 | 15 | 2 | 120 |
| 2 | 128 | 1 | 128 |

Extending the concept to two bytes (15 bit), we can look at combinations where M does not exceed 32768:

| Encoding | Range (n) | Voices (k) | Combinations (M) |
|----------|-----------|------------|------------------|
| 1 | 12 | 7 | 31824 |
| 2 | 14 | 6 | 27132 |
| 3 | 18 | 5 | 26334 |
| 4 | 28 | 4 | 31465 |
| 5 | 57 | 3 | 32509 |
| 6 | 255 | 2 | 32640 |

Tables for three, four and five bytes can be found in [17]. Even ambitious keyboard players should be satisfied with a range of 63 tones and chords of up to ten notes i.e. a five byte encoding: most semi-professional and professional keyboards have a range of five octaves (60 tones).

The five byte encoding is therefore to be considered the maximum in the current version of MINI.

Quality scaling is possible in MINI in several variations: i) switching between chord encodings with more or less bytes is supported, with the choice of limiting the tonal range or the number of voices; ii) fine-tuning MIDI messages such as “Program Change” and “After Touch” can be sorted in the order of importance for layered encoding, to be selected for inclusion or omission for quality scaling; iii) the fine-tuning messages may also be switched from “per-tone” to “per-chord” and “per-sequence”, offering another dimension of quality scaling.

Combined with a bandwidth adaptation mechanism with a smooth sending rate (in our case, M-TFRC), the frequency of changes in quality scaling was experienced to be quite tolerable in our applications and studies.

Note that the status of MINI as described represents “work in progress” in the state used for our application. Nevertheless, the authors see sufficient potential for this work to become the outset of a standardization effort.

5. GlobeMusic: a new eMusic-making scheme

Despite all the advancements of the parts introduced up to now (WorldBeat, NetMusic, RT-Score i.e. M-TFRC and MINI), the invariant problem of audible signal delays remains in a global scale application. In this respect, we believe in the need for a new cooperative eMusic-making scheme based on the experiences described in the previous chapter: WorldBeat proved the feasibility and appropriateness of new eMusic-making schemes for the individual, and NetMusic successfully adopted and adapted the cooperative eMusic-making scheme introduced by ResRocket. Note that the latter scheme is i) totally asynchronous and ii) appropriate for experienced musicians only (which is clearly not the target we had set out for application targeting the 4th mass (!) medium).

We will start this paragraph with an analogy: global videoconferencing is usually experienced as a synchronous cooperative application despite considerable delays if pairs of audio and video streams are kept lip-synchronous and if, in a dialogue, the network/system-related delay of a stream carrying a response does not exceed the usual “think-time”. In a question-answer dialogue, this constraint is usually maintained: as to timing, the human process “ask-think-reply” is usually not noticeably different from the human/network process “ask-transmit-think-reply-transmit” (the computing part is omitted here for the sake of brevity). In a heated debate, however, speakers are often interrupted while talking; such a dialogue may become noticeably different if not unfeasible in a networked videoconferencing setup since the human/network process may have considerably different timing than the human-only process.

The above analogy leads to a key observation:

Global-scale cooperation can be experienced as synchronous if the technical setup conserves sufficient fidelity of the streams (cf. lip-synchronous above) and if the scheme introduces human delays larger than (or at least in the order of) network delays.

In our case, the technical constraints are maintained with the introduction of M-TFRC and MINI. As to the scheme, the first version of GlobeMusic implements what we call *take-shift* as described below:

The scheme slightly resembles a jazz (jam session) like setup; instead of hearing himself and all others, a musician hears himself and a limited number of “others” plus a computer provided “initial” background track.

The WorldBeat “improvisation helper” metaphor is adapted; depending on the experience of the group members, the “degree of help” may vary; e.g., as described for WorldBeat and NetMusic.

For a given take, musicians are ordered, to be called #1, #2, etc. in the remainder; the lowest order #0 is assigned to the computer provided background track.

Musician #1 hears the initial background track and adds his own track with improvisation helper assistance; thereby, #1 hears #0 and #1 (background and herself).

The result is sent to #2 in real time based on M-TFRC and MINI, played out there with a slight delay (with respect to the *start* at #1, not the *end* like in ResRocket!); #2 adds her own track with improvisation helper support.

The result is sent to #3 who “joins in” a tick of a second or at most a few seconds later, and so on.

Once a take is fully recorded, it may be audited by all group members or the next take may be created.

For subsequent takes, the ordering is re-arranged, following either a cyclic shift rule, or the desires of the “band”, or a random drawing.

Videoconferencing is optionally added for a more intense group experience; a major adaptation is necessary in comparison to standard videoconferencing: during recording of a take, the usual audio channel is replaced by the MINI channel, and “chord synchrony” between the MINI and video streams (as opposed to lip synchrony) must be maintained – with track #0 defining the reference timing for all MINI tracks and for all video streams.

Fig 5.1 depicts the ensemble of GlobeMusic, including the modules introduced in the previous chapters.

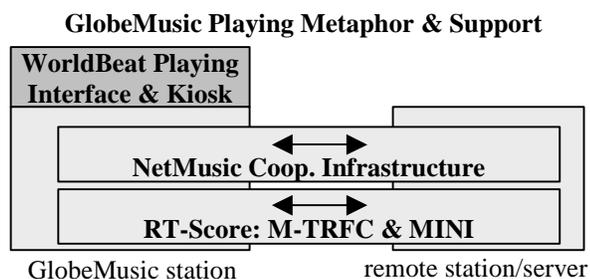


Figure 5.1: GlobeMusic Distributed SW Architecture

6. Conclusion

We introduced GlobeMusic as a combination of efforts regarding computer-human interaction, individual eMusic-making, appropriate computer network protocols, musical data formats, and cooperative eMusic-making. Major parts of the system called WorldBeat, NetMusic, and RT-Score were intensively experimented with. WorldBeat has proven the appropriateness of its approaches as an exhibit used by hundreds of thousands of users. Careful evaluation by large user groups and international review boards back this claim further. While large-scale evaluation of the last evolution step (GlobeMusic) is still to be carried out, considerable insight into the “laws” of music-making for the 4th mass medium Internet has been gained.

7. References

- [1] J. Borchers, *A Pattern Approach to Interaction Design*. John Wiley & Sons, March 2001.
- [2] J. Borchers, “WorldBeat: Designing a baton-based interface for an interactive music exhibit”, *Proc. CHI'97*, ACM 1997.
- [3] B. Braden, D. Clark et al., “Recommendations on Queue Management & Congestion Avoidance”, RFC 2309, April 1998.
- [4] M. Crovella, A. Bestavros, “Self-similarity in WWW traffic: evidence and possible causes”, *IEEE/ACM Trans. Networking* 5, 6 (Dec. 1997), pp. 835 – 846
- [5] S. Floyd, M. Handley, et al., “Equation-Based Congestion Control for Unicast Applications”, *Proc. SIGCOMM 2000*.
- [6] J. Geirland, E. Sonesh-Kedar: *Digital Babylon: How the Geeks, the Suits, and the Ponytails Fought to Bring Hollywood to the Internet*, Arcade Publ., Sept. 1999
- [7] S. Jacobs, A. Eleftheriadis, “Providing Video Services over Networks without Quality of Service Guarantees”, *WWW Consortium Workshop Real-Time Multimedia and the Web*, 1996.
- [8] B. McQueer, *The USENET MIDI Primer*, WWW site <http://www.harmony-central.com/MIDI/Doc/primer.txt>
- [9] R. Rejaie, M. Handley, D. Estrin, “An End-to-end Rate-based Congestion Control Mechanism for Realtime Streams in the Internet”, *Proc. INFOCOM 99*, 1999.
- [10] ResRocket Surfers, WWW site <http://www.resrocket.com/>
- [11] E. Scheirer, B. Vercoe, “SAOL: MPEG-4 Structured Audio Orchestra Language”, *Computer Music* 23, 2 (1999), pp 31-51.
- [12] E. Scheirer, “Structured Audio and Effects Processing in the MPEG-4 Multimedia Standard”, *Multimedia Systems* 7, 1 (Jan 1999), pp 11-22.
- [13] E. Selfridge-Field (Ed.), *Beyond MIDI*, MIT Press 1997.
- [14] D. Sisalem, H. Schulzrinne, “The Loss-Delay Based Adjustment Algorithm.”, *Proc. NOSSDAV'98*, 1998.
- [15] V. Jacobson, “Congestion Avoidance and Control”, *Proc. SIGCOMM 1988*, pp. 314-329.
- [16] J. Webers, *AudioStudios (German)*, Poing: Franzis, 1994
- [17] M. Welzl, *NetMusic: Real time concepts and systems for telecooperative exchange of musical informatik (German)*. Diploma Thesis, University of Linz: Telecooperation, 1997.
- [18] www.ietf.org/internet-drafts/draft-lazzaro-avt-mwpp-midi-mpeg4-00.txt
- [19] E. Zugra, K. Calvert, S., Bhattacharjee, “How to Model an Internetwork”, *Proc. IEEE Infocom '96*, San Francisco, CA.