



**CONSIGLIO NAZIONALE
DELLE RICERCHE**

Goffredo Haus [^], Alberto Sametti [^]

**Scoresynth: A System for the Synthesis of Music
Scores Based on Petri Nets and a Music Algebra**

N. 7/89

8/92

**PROGETTO FINALIZZATO
SISTEMI INFORMATICI E CALCOLO PARALLELO**

SOTTOPROGETTO 7

SISTEMI DI SUPPORTO AL LAVORO INTELLETTUALE

**PROGETTO FINALIZZATO
SISTEMI INFORMATICI E CALCOLO PARALLELO**

SOTTOPROGETTO 7
Sistemi di Supporto al Lavoro Intellettuale
Coordinatore Cesare Maioli

Goffredo Haus [^], Alberto Sametti [^]

Scoresynth: A System for the Synthesis of Music
Scores Based on Petri Nets and a Music Algebra

N. 7/89

8/92

Reprint

[^] LIM - Laboratorio di Informatica Musicale
Dipartimento di Scienze dell'Informazione
Università degli Studi di Milano.



ScoreSynth: A System for the Synthesis of Music Scores Based on Petri Nets and a Music Algebra

Goffredo Haus and Alberto Sametti

Introduction

In this paper we show how music structures can be described, processed, and synthesized by means of a more flexible kind of representation than the staff. In fact, symbols are chosen and organized depending deeply on instrumental needs within common music notation. The level of representation within scores is more detailed than that of music composition and more abstract than that of timbre modeling. The new kind of representation we are proposing makes up a conceptual music framework with as many different levels of abstraction as the musician needs and allows us to explicitly describe and process what we call *music objects* (MOs) — both traditional music and nontraditional music objects. An MO means anything that could have a musical meaning and that we think of as an entity, either simple or complex, either abstract or detailed — an entity with a name and some relationship with other MOs.

Therefore, we can describe MOs at various abstraction levels within a hierarchical context of description. For example, we can have the following: the structural level, the score level, the timbre level, and any other level we think is suitable for our music representation. Common music notation is characterized by many different languages (one or more languages for each level of representation). Furthermore, modern composers have introduced a myriad of notation conventions to represent concepts that the common music notation cannot represent. Our work is devoted to the definition of a *one and only language*, suitable both for every level of representation of common music information and for contemporary music representation (i.e., *concrete*, *electronic*, and *computer music*).

To identify the most suitable description tool, we have sought a formal tool that

- requires few symbols,
- has a graphical form of notation,
- allows hierarchical description,
- allows the description of MO processing,
- allows time description,
- allows deterministic/nondeterministic models,
- allows macro definitions for common structures, and
- shows the musician the score synthesized as the model execution goes on.

Score modeling by petri nets and a music algebra

Since 1980 we have been using petri nets (PNs) as the basic tool for music description and processing.¹ Other experiments have been carried out by Pope² based on predicate-transition nets.

Throughout our research, we have developed some programs to edit and execute PNs for both music and general-purpose applications. The most recent program we have developed is ScoreSynth, where we use PNs for describing, processing, and synthesizing music scores. While we associate MOs to place nodes for describing music information, we associate transformation rules to PN structures and PN parameters (marking, numeric labels, algorithms, etc.). MOs may be processed modifying the superimposition and juxtaposition laws within PN structure. On the other hand, PN parameters allow us to create *instances* of MOs which modify themselves according to the behavior of PN models during executions. Furthermore, we can represent the same musical information with PNs at a lower or higher level of abstraction by means of suitable alternative modeling approaches. The modeling of music with PNs makes natural the description of the *structural* level of music as a multilevel environment within which MOs may flow concurrently and interactively. The more detailed levels of representation (the score and timbre levels) may be described and/or processed depending on user choices.

Indeed, a PN model with a particular initial marking may represent a family of scores; a particular execution of the model synthesizes a specific score. By changing the initial marking of the model, we change the family of scores that can be produced by model executions. In a special case when we have a fully deterministic model, we can produce only one score (given a particular initial marking).

The following paragraphs discuss the definition and implementation of music objects, petri nets, and music algorithms within the ScoreSynth system. ScoreSynth has been developed on the Macintosh family of computers within the Macintosh Programmer's Workshop (MPW) development environment.

Music objects

The meaning of MO, as already mentioned, is strictly related to the place nodes in our implementation. By MO we mean not only any sequence of notes but also something more general: something that is not exclusively linked to listening and is not connected with the idea of process, because it could last zero seconds, as far as the measurement of time is concerned. For example, an MO may be a control command for a sound synthesis device. Obviously, the definition of MO is partly affected by the standard of the music code we use (MIDI) and must stay within the bounds of this standard. Appendix A gives the complete Backus-Naur form definition of MOs.

In this section we go through all kinds of MOs. To simplify the explanation, we name, without distinction, a place or its associated MO everytime it is obvious. The reader must keep in mind, however, that a placenode and an MO are two very different things. Moreover, we want to remember that listening to a score happens in a different time than its synthesis, that is, the performance of the PN models.

Notes. If an MO is formed only by notes (see Figure 1), each note must be specified through the parameters necessary for its reproduction through MIDI sound generators. These parameters are the MIDI channel (that we can nearly compare with the instrument), the pitch, the intensity, and the duration. The duration is expressed implicitly, by giving the starting and the ending point of a note. We want to remember that the MIDI standard imposes some limitation on the values of these parameters, which we discuss in another section. Literature on the MIDI standard is widely available.

Metaevents. All the events inserted in the coding of MOs that give general information on the score but are not directly performed because they cannot be coded by MIDI codes belong to the metaevent class.

Metaevents include the tempo (beats per minute), the base (ticks per quarter note), the time



Figure 1. A simple music object (MO).

signature (4/4, 3/4, 6/8, etc.), and rests. If there are rests, they must form a separate single MO because they cannot be explicitly described by MIDI codes.

The general lines of the work impose an intrinsic limitation on the score synthesized by the execution of a net. The first three types of metaevents are not coded directly into the score, but they are only a mark of the last modification. All MIDI events, which build up the score, have an absolute timing reference expressed in multiples of 0.5 ms. For this reason, it is not possible to reconstruct a score in traditional notation if there is just one modification, for example, in the time signature. This problem does not occur when the tempo parameter remains fixed from beginning to end. This limitation affects the conversion of a PN model into a traditional score, but not the performance of the model via MIDI.

Other MIDI messages. All MIDI messages not directly concerned with a note can be put in an MO. If the message belongs to the class *channel voice*, it can be put in the note event; if the messages belong to other classes (i.e., *channel mode*, *system common*, *realtime*, *exclusive*), the MO must be formed by only one of these messages and nothing else.

Coding formats. MOs can be coded according to three syntaxes: a coding language unique to ScoreSynth, a portable alphanumeric MIDI coding language, and a coding format that implements the Standard MIDI Files (SMF) 1.0 format (both type 0 and 1). The first format is used internally by the program. The second format has been chosen for the sake of editing tool availability (i.e., common word processors). The third format has been implemented for file import/export. So MOs may be defined either by a common text editor or a MIDI controller (a keyboard, a guitar, a microphone, etc.).

Before executing a PN model, ScoreSynth converts all the MOs (both ASCII and SMF) into the inner MIDI format, the only format that ScoreSynth directly works with. If the original files don't undergo modifications, they can be expelled to make the following executions quicker.

Building a net, the user can decide if a place node must have only a control role, namely, without any chance to have MOs associated. On the other hand, the user can choose if the MO must be always saved on a file (permanent MO) for possible uses at the end of the execution, or kept in the memory (volatile MO), and so remain only till the end of the execution. These choices cannot be modified during the execution.

There are two other possibilities for the user. The first is to redirect all the MIDI events of an MO on a single specified channel among the attributes of the place to which the MO is associated. The second is to associate an MO to a place and to set whether it must be executed or it must be subject to following transformations.

Petri net implementation

A petri net is usually defined by a quadruple $\langle P, T, A, M \rangle$, where P is the set of places, T is the set of transitions, A is the set of arcs connecting both places to transitions and transitions to places, and M is the marking of the net. Our implementation of petri nets extends the "classic" definition of Petri³ (which is well exemplified with respect to modeling features of PNs in Peterson⁴), so that the




Figure 2. Petri net icons.

musician can build up hierarchical music models defining MOs and music algorithms (MAs) made up of operator expressions. In this way, a petri net is defined by means of the triplet $\langle P, T, A \rangle$ and the following three statements:

- (1) $P \ni p \equiv \langle \text{identifier}, \text{tokens}, \text{capacity}, \text{MIDI channel}, \text{object}, \text{play}, \text{file} \rangle$
- (2) $T \ni t \equiv \langle \text{identifier}, \text{algorithm} \rangle$
- (3) $A \ni a \equiv \langle \text{node-from}, \text{node-to}, \text{multiplicity} \rangle$

Here P is the set of places, T is the set of transitions, and A is the set of arcs. Also, *identifier* is the label of the node (both place and transition); *tokens* is the number of tokens within the place, that is, the value of the marking function M (see below); *capacity* is the maximum number of tokens allowed for that place; *MIDI channel* is the number of the MIDI channel where the MO, if one exists, has to be played; *object* defines whether an MO may be associated to the place or not; *play* defines whether the object has to be played or not; and *file* defines whether the MO associated to the place, if one exists, is stored on file or not. In the second statement, *algorithm* defines whether the MA associated to the transition, if one exists, has to be executed or not. In the third statement, *node-from* is the identifier of the node that starts the arc, *node-to* is the identifier of the node that ends the arc, and *multiplicity* is the numeric label of the arc. If a place has an associated MO stored on file, its *identifier* represents both the label of the node and the file identifier.

P , T , and A are finite sets; P cannot be empty. Both objects and algorithms have their own syntax definitions; they are discussed respectively in Appendixes A and B to this paper.

The marking of a net is the distribution of tokens into the place nodes; their number and distribution may change during net execution. A marking function is defined as the following array:

$$M \equiv (m_1, m_2, \dots, m_i, \dots, m_n)$$

where n is the number of places within the net, and N is the set of positive integers and every $N \ni m_i = M(p_i)$ with $P \ni p_i$.

Then, the hierarchy of the model is described by the set

$$S \ni s \equiv \langle \text{identifier}, \text{node-begin}, \text{node-end}, \text{recursion-level} \rangle$$

which is the set of hierarchy links, where *identifier* is the label of the net, *node-begin* is the identifier of the beginning node for the morphism application, *node-end* is the identifier of the ending node for the morphism application, and *recursion-level* is the level of recursion allowed for that net. We discuss how hierarchy and recursion are implemented in the ScoreSynth program in the sections entitled "Morphisms" and "Recursion."

The S set represents all the information about the hierarchy of the model, while the P , T , and A




Figure 3. (a) Firing inhibition due to the capacity of the output place. (b) Firing conflict due to the capacity of the output place.

sets completely describe a particular net of the model. In other words, a model is fully defined by an S set and as many triplets $\langle P, T, A \rangle$ as the number of the nets within the model.

Petri nets are commonly represented in graphic form. We too use a graphic form of representation with the set of icons shown in Figure 2.

The icons represent (from left to right) a named simple place with tokens and capacity attributes, a named place that represents a subnet (see the section on morphisms), a horizontal named simple transition, a vertical named simple transition, a horizontal named transition that represents a subnet, a vertical named transition that represents a subnet, and an arc with the multiplicity attribute set to 4.

The firing rule. The execution of a net is made up of transition firings. A transition may fire if each one of the input places — that is, places connected with oriented arcs to the transition — has at least one token. The transition firing has two effects: to decrement the marking of each input place by one token and to increment the marking of each output place (i.e., a place connected with an oriented arc from the transition) by one token.

In the following paragraphs we see how this basic firing rule is affected by capacity, multiplicity, and timing parameters. This firing rule can be considered as a particular case of the firing rule we have defined for our extended implementation of PNs. From this point of view, we can describe the basic firing rule with the following settings: All multiplicities are equal to 1, all capacities are infinite, and all durations of transition firings are null.

After the starting of a net, firings follow one another until there are no more transitions which may fire. At the end of transition firings, the execution of the net stops.

Capacity. Capacity is a feature of place nodes which defines the maximum number of tokens allowed for each place. The capacity feature implies that the firing rule previously defined has to be modified. Conditions of transition firings do not change with respect to the input places, while transitions cannot fire if the marking of one output place, at least, will exceed its capacity after transition firing; that is, the marking of that place added by the number of tokens due to transition firing is greater than its capacity.

Two simple examples show this concept. In Figure 3a transition $t1$ cannot fire because place $p2$ is marked with as many tokens as its capacity. In Figure 3b only one transition can fire, either $t1$ or $t2$, but not both. This simple net is a conflict structure, which is nondeterministic.

Multiplicity. To improve petri nets' descriptive power and promote graphic simplicity, we have assumed the multiplicity extension, that is, the ability to add numerical labels on the arcs. The multiplicity extension is a partial implementation of self-modifying nets, which were introduced by Valk.⁵

Also, the multiplicity feature requires modifications to the firing rule: A transition may fire if each one of the input places has at least as many tokens as the numerical label on the arc connecting the




Figure 4. An example of transition firing: (a) transition t1 may fire; (b) the effect of the firing of transition t1.

place to the transition. The transition firing now has the following effects: it decrements the marking of each input place by as many tokens as the numerical label on the arc connecting the place to the transition, and increments the marking of each output place by as many tokens as the numerical label on the arc connecting the transition to the place. Figure 4 shows the behavior of a transition firing within a simple net with multiplicity on some arcs.

Timing

PNs are particularly suitable for describing concurrent processes and for controlling their synchronization or lack thereof. When a transition may fire, we don't know when it fires, and the duration of firing is assumed to be instantaneous. When we use PN to describe asynchronous models, we have no problem with this feature. There are no measurements of time within petri nets: It is the behavior of the net that implicitly determines timings. It is the structure of the net that determines the firing sequence, without respect for their durations. On the other hand, we can synchronize MOs simply by the suitable structuring of the nodes' connections.

Finally, the firing sequence changes among different executions of the net, so there is no correspondence between a net and its firing sequence. We can have many transitions qualified for firing at the same time, and we cannot know which one fires first: The net does not describe this kind of information. So every execution of a PN model may give a different firing sequence.

This feature may seem to be too abstract for the modeling of real complex systems. Several approaches are known regarding how to describe timed events' durations within PN. Most of them represent timed events by means of timed transitions. Our approach associates timed MOs to places and their transformations to transitions. Both the firing of a transition and the execution of an associated MA have a null duration. In this way, when a token is put into a place with an associated MO, the token cannot be considered for the firing of transitions connected to the place until the associated MO has ended.

Consider the net in Figure 5. There are two places with associated MOs: MO A with place A and MO B with place B. This example is concerned with the duration of processes but not with their

meaning. MO A lasts 1 second and MO B lasts 6 seconds. Every token put into place A is not disposable for 1 second; tokens put into places p1 and p2 are immediately disposable; that is, the amount of time between the firings of t1 and t2 is null.

If place A has one token, then t2 cannot fire due to the capacity of place A; if place B has one token, then t3 can add tokens to place B (six is the maximum) when it fires. The effect is that a certain number of instances of MO B may run concurrently.

Let us analyze some steps of the net execution. Figure 6 shows the marking at time 0 after both t1 and t2 firings; we can see that both places A and B have their own token, so that MO A and MO B are running (see Figure 7). At the end of MO A (the shorter one), the token of place A is




Figure 5. Example of timed net.




Figure 6. The marking at time 0.




Figure 7. MOs started at time 0, after t1 and t2 firings.

disposable and starts the firing of **t3**, which starts the second instance of **MO B** running. At time 1, one of the tokens of place **p2** lets **t2** fire, which starts **MO A** running (see Figures 8 and 9).

The firing sequence happens every second in the same way four times. The last firing happens at time 5; the complete result of the net execution is shown in Figure 10.

Morphisms

Theoretical literature about morphisms and related formalisms is broad and multifaceted. We limit our discussion to a particular kind of morphism that seems to be suitable for our purposes: the




Figure 8. The marking at time 1, after both **t3** and **t2** firings.




Figure 9. Running MOs at time 1; new MOs started at time 1.




Figure 10. The complete map of the MOs generated at the end of the firing sequence.




Figure 11: (a) The net we have designed. (b) The high-level net. (c) The p2 subnet.

refinement morphisms. The hierarchical approach introduced by Kotov⁶ is rather close to our implementation. Refinements can define very complex PN models by means of simple PNs and hierarchical structures — that is, allowing models to be designed by either a top-down or a bottom-up approach. A refinement that we call a *subnet* is a PN which gives a more detailed description of a node (either a place or a transition) of the upper level of abstraction. We call that node the “father” node and that subnet the “daughter” net.

If we choose to define a subnet associated to a place, then the daughter net must have two special places: the input place and the output place. Input arcs of the father place are input arcs of the input place of the daughter net; output arcs of the father place are output arcs of the output place of the daughter net. Transitions can be refined in the same way. An example of a subnet which refines a place is shown in Figures 11a, 11b, and 11c.

The conceptual framework of our approach is discussed in Haus and Rodriguez.⁷ This approach describes a whole complex model as a single large net. This large net may be either the result of many refinement steps in the frame of a top-down detailing (see also Goguen⁸) or the starting information basis on which we can apply abstractions as many times as we want to implement the hierarchical structure in the frame of a bottom-up modeling (see also Bertoni et al.⁹).

Recursion

Recursive nets can also be modeled. Both direct and mutual recursion are allowed. In ScoreSynth to avoid recursive expansions with no end, we must assign to every recursive net its recursive level number. That is to say, we have to specify how many times the net can call itself. Figure 12 shows an example. In this model we use a recursive net: **p1** (Figure 12b). If we choose a recursion level equal




Figure 12. (a) Simple recursion example net. (b) Recursive subnet p1.




Figure 13. The result of the expansion of place p1 with recursion level 1.



Figure 14. The two MOs p1 and p2.




Figure 15. Elementary PN structures.




Figure 16. Two different PNs (a, b + c) for the loop structure and the score resulting from their execution.

to 1 (only one call of a net inside itself), we obtain the expanded net of Figure 13, in which the dashed lines represent the deleted arcs during the sequence of recursive refinement steps (depth first) that give the expansion, and the bold lines represent the added arcs needed to implement refinement morphisms. Input is the input place of net **p1**. **p1** is the output place of net **p1**. Numbers in the small squares near the bold lines indicate the step of the expansion in which the arcs are added. Numbers in the small circles near the dashed lines indicate the step of the expansion in which the arcs are deleted.

Structuring music objects

Starting from two simple MOs (**p1** and **p2** of Figure 14), we can see the elementary structures in Figure 15 necessary to design more complex nets. The places with a name are those to which an MO is associated. When a token reaches a place with an MO, it activates a "process" that is the execution of the MO itself. This "process" stays in the *running* state for all the duration of the MO. In Figure 15 we also show the scores resulting from the executions of the corresponding nets.

A very special case is the third one. This kind of net introduces nondeterminism. In this case, ScoreSynth determines the firing sequence with a pseudorandom uniform distribution. All the considered elementary structures can freely be joined to form more complicated PN structures.

A little more complex structure is the loop structure. Figure 16 shows two possible implementations.

Music algorithms

In this paragraph we briefly discuss how we have implemented the ability to work with MOs in ScoreSynth. The main idea is to allow transformations of both the parameters of sound (pitch, timbre, intensity, and duration) and the order in which notes appear. Appendix B gives the complete Backus-Naur form definition of music algorithms (MAs).

MAs join with transitions. If a transition has an associated MA, when the transition fires the MA is applied to the MOs associated to the input places and the outcoming object is placed in all the

output places allowed for MOs. During the processing of an MA, it is also possible to use information from every MO available at that moment. We have observed that the modeling of music is as greatly improved as the expressive power of the formal language adopted if it is possible to have all the model information at our disposal. We have studied formalisms that allow us to reach all nodes within the model without loss of time. In early research, we have considered an approach similar to that of the statecharts introduced by Harel et al.¹⁰ Then, we have defined a simpler mechanism that allows the immediate availability of all the information of the whole PN model as possible input arguments for MAs.

An MA can be formed by one or more single transformations. Each transformation may affect either the notes' parameters or their order. Before making a transition fire, ScoreSynth merges all the MOs in the entrance of a transition. It applies the transformation on the single MO obtained and associates the outcome of the transformation to each of the output places to which it is possible to associate an MO. The transformation can be applied only to note events. All the other MIDI commands are filtered and lost.

Metacharacters, constants, and variables. If *X* stands for the object on which the MA is processed at the input, the MA can be applied to the whole *X* or just to one of its subsequences, chosen by the user by the definition of the positions of both the first and the last note. The MA is applied to the range of notes which is set in its header, to each note as it occurs. The header also defines the kind of parameter to which the operator expression must be applied: pitch, duration, intensity (key velocity), MIDI channel, and the order of the affected notes.

When we draw up an operator expression, we can use metacharacters — that is, special characters — which are needed to determine relevant transformations. The implemented metacharacters are \$, ?, !, and %. While \$ and % are the same as constants, ? and ! are variable quantities. We examine them one by one:

\$ contains the total number of notes in *X*;

% contains the total number of notes in the subsequence of *X* on which we have defined to operate;

? contains the value of the parameter, referred to the current note, that we want to change; and

! is an index that contains the position of the note on which the operator expression works; the position value is referred to the first note to be changed within the sequence.

Operators' syntax. As we have already mentioned, the parameters of the sound that we can modify at the level of MIDI code are pitch, duration, MIDI channel, and intensity (key velocity).

Without dwelling on the concept of channel in MIDI, we want to explain that a channel stands approximately for an available instrument. We cannot directly operate on the characteristics of the sound, but only change the channel on which we want to send the notes.

The nucleus of an operator expression is simply an integer expression, that is, an expression with operators +(sum), -(subtraction), *(multiplication), /(division), and parentheses (). The terms of an operator expression can be integers, metacharacters, and alphanumeric pitch codes.

Here is an example to show a syntactically correct operator expression, with symbols, pitch, metacharacters, and integers:

$(2 * C3 / 4 - ?) * !$

Each of these operator expressions must be preceded by a header that specifies to which parameter the operator expression is applied and which part of an MO must be transformed.

We mention the meaning of each terminal character that identifies the kind of transformation:




Figure 17. A simple net with an MA.

C (channel)	operator on channels
D (duration)	operator on durations
L (loudness)	operator on intensities
M (multiply)	operator for multiplying notes
R (rotate)	operator for rotating notes
P (pitch)	operator on pitches
I (inversion)	operator for inverting the ordering of notes
K (kill)	operator for deleting notes
S (save)	operator for preserving notes

To use information about any MO with respect to duration (rhythm), pitch, MIDI channel, and intensity, it is possible to specify the identifier of the MO and the starting note inside it.

For example, **P: 1, \$, [Theme, 1], ? + 2** means that, to compute the new values for the pitches of the current MO, the pitch values of the notes within the MO **Theme** (starting from the first note) are used. In this example, metacharacter **?** refers to the MO **Theme**, while **\$** refers to the current MO. See Figure 17 for a simple net that shows the use of such an MA associated to the **Alg** transition. The new MO obtained from the transformation of **Theme** is joined to the place **Theme.mod**.

If the upper bound of the range exceeds the ordering number of the last note in **Theme**, then it is automatically modified to fall inside the bounds. If the current MO is empty, then not only the pitch values of **Theme** are taken, but also channel, duration, and intensity values, according to the range specifications.

Operators' examples. Starting from an MO, Figure 18 shows some examples of feasible operators. On the left is the operator and on the right the resulting MO.

Operators on pitches can even be applied according to scales different from the chromatic one. This must be done if we want the outcoming sequence in the same tonality as the beginning one. ScoreSynth has the following set of scales on which we can operate: *major*, *harmonic minor*, *pentaphonic*, *minor thirds*, and Debussy's *whole tones*. The third operator, for example, uses a C-major scale as indicated by the statement **[C]**.

These are the choices made by ScoreSynth when the result of operator expressions goes beyond upper and lower bounds:

- *Channel* (1..16). We obtain the channel with a mod 16 operation.
- *Pitch* (0..127, C -2..G 8). If the value obtained is more than 127, then we subtract the cardinality degree of the current scale as many times as is needed to obtain an allowable pitch value. If the value obtained is less than 0, then we must add the cardinality degree of the current scale as many times as necessary to obtain an allowable pitch value.
- *Duration*. If it is less than 0, then it is set to 0.
- *Intensity*. If it is less than 0, then it is set to 0. If it is more than 127, then it is set to 127.

		{ The starting MO }
P: 1, \$, ? + 1		{ Tonal transposition }
P: \$ - 2, \$, 2 * G3 - ?		{ Specular inversion }
P [C]: 1, \$, ? + 1		{ Degree transposition }
D: 1, \$, ? * 2		
L: 1, \$, ! * (127 / %)		{ Crescendo }
L: 1, \$, (% - ! + 1) * (127 / %)		{ Decrescendo }
I: 2, 5		{ Retrogradation }
R: 2, 5, 1		{ Rotation }
M: 2, 5, 2		
K: 2, 6		
S: 5, 7		

Figure 18. Examples of operators.

"Macro" petri nets

When we are designing a net model, we often have to create many similar nets, which differ in their place attributes and MAs, but which are identical in their structures. With a *macro* net we can use just one net as the base model and then modify its attributes and/or MAs at our need. So it is also possible for the user to create his own libraries, containing commonly used PN models. Appendix C gives the complete Backus-Naur form definition of macro nets.

The way we can do this is by writing, for every subnet place or subnet transition we want, a modifier list for the net, which is loaded into memory and actualized before the execution of the whole expanded net.

An example that shows the use of a macro net invoked by a modifier list is included in the net model of Figure 19. We choose to expand place p2 using the following modifier list:

Input: Theme	{ change the name of place "Input" and tell not to play the music object associated to the place theme }
Output: C = 2	{ change the capacity of place "Output" }
t1: Alg (C: 1, \$, 3	{ change the name of transition "t1" and add an algorithm consisting of two operators }
P: 1, \$, ? + 1)	
Alg -> Output = 2	{ change the multiplicity of arc "Alg -> Output" }

Score synthesis

Model execution

Here we briefly discuss the main steps of the model execution algorithm. Appendix D gives a pseudocode version of the algorithm, but first we introduce such a "code" with a natural language explanation.

The main loop is included in the *Execute* procedure. The *Execute* procedure controls again and again the passing from a temporary blocking to the following firings; a temporary blocking happens when all the tokens within a net are in places either with an associated MO or without output arcs or with all output transitions that cannot fire.

The minimum time interval (*DeltaTime*) is computed for every cycle among all the active MOs within which one process at least terminates. Next, all active processes are executed for a *DeltaTime* number of time units. The variable that represents the virtual clock is increased by *DeltaTime* units, and all terminated processes are deleted from the list of active processes. Execution stops when the *Firings* function evaluates that no transitions have fired and no active processes exist.

The *Firings* function takes care of transition firings. It simulates again and again the execution of all transitions either which may fire in the frame of alternative and/or conflict net structures or which may have simple firings (i.e., which don't belong to alternative and/or conflict net structures), until there are no more transitions that may fire. The *Firings* function returns control to the *Execute* procedure when there is a temporary block; that is, no more transitions may fire at that time.




Figure 19. (a) The top-level net. (b) The subnet p2. (c) The net resulting from the macro expansion.

The ScoreSynth Executor

In the Executor environment the user can interact with some window dialogues that we are going to describe. Figure 20 shows the main dialogue. It gives to the user some facilities concerning the playback of the generated scores. The user can control up to 10 scores in this dialogue and can modify some listening parameters of the current selected score. It was not our intent to design a complete sequencer environment. If the user wants all the features included in any commercial sequencer, he has only to export the score as a Standard MIDI File.

These are our sequencer features: play, stop, pause, and rewind the song; change the tempo velocity in the column named Vel; start playing the score at different points using the scale under the pause button; loop the song from the starting point to the end; and reassign MIDI channels defined during the PN model execution phase to other channels.

It is also possible to interact with the execution of the net model invoking the music objects map dialogue that we'll discuss in the next paragraph.

Music objects map

A musician trying to create a net model suitable for his purposes may often want to see on the screen how the execution of his model proceeds and the causal relationships between transitions' firings and score synthesis, that is, MOs appending in the score. In this way, he can stop the execution as soon as he finds a bug in the model and go to the editor environment to modify it. This is possible simply by running the execution of the model and looking at a graphic window that shows on a time-MIDI-channel plane the map of MOs as they are appended in the score.

Two possibilities are available: either simply run the execution and look at the growing map, or run the execution, while looking at the map, and automatically stop it every time a new MO is created (i.e., every time a token reaches a place with an associated MO).

In the last case, the user who discovers a bug (e.g., a token gone elsewhere) can obtain all the information about places' and transitions' attributes and make the needed improvements on the current model.


Score		Vel	Tape	☐ Loop											
			Rewind  Stop  Play  Pause 												
															
MIDI Channel Map															
1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16

Figure 20. The Executor window.

Figure 21 shows the dialogue that displays in the upper part the running execution in a graphic form. The vertical scroll bar lets the user move on the MIDI channel axis (1 to 16); the horizontal scroll bar lets the user move on the time axis. The user can select the time scale opening a dialogue by clicking the "scale" box near the point (0, 0). In the lower part, the dialogue shows all the information the user may need to control the execution of the net model (all the attribute values of places and transitions every time the user stops the execution).

Ending remarks

Our research has shown that music structures within PN models can be transformed simply by

- modifying markings, MIDI channel specifications, capacities of places;
- modifying labels of arcs;
- modifying MOs associated to places;
- modifying MAs associated to transitions;
- modifying the structure of PNs; and
- executing nondeterministic PN models many times.

While the editing of PN structures and parameters affects causal/structural relationships within the architecture of scores, the editing of both MOs and MAs affects the basic information units and their transformations. The ScoreSynth system seems to be a powerful means for describing and processing music, closer to music thinking and perception than common music notation. People interested in listening to pieces synthesized by PN model execution can consider the collection "Musiche per poche parti" by G. Haus (compact disk Stile Libero/Virgin Dischi, SLCD1012, Milan, 1990). Other music productions are currently in work.




Figure 21. Object map dialogue with attributes.

Acknowledgments

Many thanks to E. Bianchi, A. Camurri, G. Degli Antoni, G. Jacomini, A. Rodriguez, and R. Zaccaria for their contributions to previous research in and experience with petri nets and music. Particular thanks to S. Scolaro for his relevant contribution to the design and development of some features of the ScoreSynth Release 2.0 (graphic interface and PN macros, especially). Moreover, special thanks are due to the LIM (Laboratorio di Informatica Musicale) staff for the scientific and technical support provided throughout our research.

This research is partially supported by the Italian National Research Council in the frame of the Music Topic (LRC C4, grants No. 89.00031.69, 90.00678.PF69, and 91.00904.PF69): "Intelligent music workstation," Finalized Project: Sistemi informatici e calcolo parallelo.

Appendix A: Music objects Backus-Naur forms

The following Backus-Naur forms describe the syntax of MOs:

< MusicObject >	::= < MusicObject > . < MetaEvent > < MusicObject > < MIDIEvent > .< MetaEvent > < MIDIEvent >
< MetaEvent >	::= Tempo < integer > Base < integer > Signature < integer > < integer > Rest < integer > MidiEx < MIDIExBytes >
< MIDIExBytes >	::= "...any sequence of bytes, in hex, according to MIDI..."
< MIDIEvent >	::= < Header > Pgm < integer > < Header > After < integer > < Header > Polyp < integer > < Header > Ctrl < integer > < integer > < Header > Bend < integer > < Header > < NoteEvent >
< Header >	::= < Channel > < Quarter > < Offset >
< NoteEvent >	::= < Duration > < Pitch > < KeyVel >
< Pitch >	::= < Note+Mod > < Octave > < KeyNumber >
< Note+Mod >	::= < Note > < Note > #
< Note >	::= A B C D E F G
< Channel >	::= 1 ... 16
< Quarter >	::= < positive integer > { ...depends on song length }
< Offset >	::= < positive integer > { ...depends on .Base }
< Duration >	::= < positive integer >
< KeyVel >	::= 0 ... 127
< Octave >	::= -2 ... 8
< KeyNumber >	::= 0 ... 127

Appendix B: Music algorithms Backus-Naur forms

The following Backus-Naur forms describe the syntax of MAs:

< Algorithm >	::= < Algorithm > < Operator > < Operator >
< Operator >	::= < QualOp1 > : < RangeExpression > , < RangeExpression > < QualOp2 > : < RangeExpression > , < RangeExpression > , < OpExpression > < QualOp3 > < TonalExpression > : < RangeExpression > , < RangeExpression > , < OpExpression > < QualOp3 > < TonalExpression > : < RangeExpression > , < RangeExpression > , [< MusicObjectId > , < RangeExpression >], < OpExpression > < QualOp4 > : < RangeExpression > , < RangeExpression > , < OpExpression > < QualOp4 > : < RangeExpression > , < RangeExpression > , [< MusicObjectId > , < RangeExpression >], < OpExpression >
< QualOp1 >	::= <i>I</i> <i>K</i> <i>S</i>
< QualOp2 >	::= <i>M</i> <i>R</i>
< QualOp3 >	::= <i>P</i>
< QualOp4 >	::= <i>C</i> <i>D</i> <i>L</i>
< MusicObjectId >	::= "... alphanumeric string ..."
< TonalExpression >	::= [< Note > < Tonality >] [< Note >]
< Tonality >	::= <i>m</i> - <i>t</i> <i>p</i>
< RangeExpression >	::= < RangeExpression > + < RangeTerm > < RangeExpression > - < RangeTerm > < RangeTerm >
< OpExpression >	::= < OpExpression > + < OpTerm > < OpExpression > - < OpTerm > < OpTerm >
< RangeTerm >	::= < RangeTerm > * < RangeFactor > < RangeTerm > / < RangeFactor > < RangeFactor >
< OpTerm >	::= < OpTerm > * < OpFactor > < OpTerm > / < OpFactor > < OpFactor >
< RangeFactor >	::= < integer > < MetaChar > (< RangeExpression >)
< OpFactor >	::= < Pitch > < integer > < MetaChar > (< OpExpression >)
< MetaChar >	::= \$! % ?
< Pitch >	::= < Note+Mod > < Octave > < KeyNumber >
< Note+Mod >	::= < Note > < Note > #
< Note >	::= <i>A</i> <i>B</i> <i>C</i> <i>D</i> <i>E</i> <i>F</i> <i>G</i>
< Octave >	::= -2 ... 8
< KeyNumber >	::= 0 ... 127

The three integers that form the header specify the MIDI channel to which the event and the instant when the event must take place are referred. The first stands for the MIDI channel. The second and the third show the moment connected with the *Base* and the *Tempo* chosen.

For example, if we have chosen *Base* 24, a header equal to 1 2 12 shows that an event must be sent on channel 1 in the middle of the second quarter, namely, after a quarter plus an eighth. In the example at the end of this paper, we give a complete MO coded with respect to this syntax.

These are the choices made by ScoreSynth when the values of the *Range* go beyond the upper and lower limits: If the lower limit is less than 1, then it is set to 1; if the upper limit is greater than the maximum number of notes in the MO, then it is set to the maximum value allowed.

Appendix C: Macro nets Backus-Naur forms

The Backus-Naur forms listed below define the syntax of modifier lists:

< ModifierList >	::= < ModifierList > < Modifier > < Modifier >
< Modifier >	::= < PlaceModifier > < TransModifier > < ArcModifier >
< PlaceModifier >	::= < PlaceId > : < PlaceOpList >
< PlaceOpList >	::= < PlaceOpList > < PlaceOp > < Place Op >
< PlaceOp >	::= < PlaceAttribute > < PlaceSetting >
< PlaceAttribute >	::= <i>play</i> <i>mute</i> <i>file</i> <i>notfile</i> <i>obj</i> <i>notobj</i> <i>simple</i> <i>subnet</i>
< PlaceSetting >	::= < NewId >
	<i>C</i> = < Capacity >
	<i>M</i> = < Token >
	<i>CH</i> = < Channel >
	/ < MusicObject > /
< TransModifier >	::= < TransId > : < TransOpList >
< TransOpList >	::= < TransOpList > < TransOp > < TransOp >
< TransOp >	::= < TransAttribute > < TransSetting >
< TransAttribute >	::= <i>enabled</i> <i>disabled</i>
< TransSetting >	::= < NewId >
	/ < Algorithm > /
< ArcModifier >	::= < PTArcOp > < TPArcOp >
< PTArcOp >	::= < PlaceId > -> < TransId > = < Multiplicity >
< TPArcOp >	::= < TransId > -> < PlaceId > = < Multiplicity >
< Multiplicity >	::= <i>1</i> ... <i>240</i>

Appendix D: The model execution algorithm

The algorithm listed below is the most abstract level of the PN execution algorithm we have defined; it is written in Pascal-like code:

```
procedure Execute;
begin
  repeat
    Blocked := NOT ( Firings ); {see the Firings function below}
    if ... no music process is active... then No_Active_Places := TRUE
    else
      begin
        No_Active_Places:= FALSE;
        compute DeltaTime for the processes (one or more) that terminate as the first;
        execute all active processes for a DeltaTime number of time units;
        delete all terminated processes from the list of active processes;
        increase of DeltaTime units the virtual clock variable;
      end;
    until Blocked & No_Active_Places;
end;

function Firings: boolean; {It returns TRUE if at least one transition fires}
begin
  repeat
    create the list of transitions which may have simple firings;
    create the list of transitions which may fire in the frame
    of alternative and/or conflict net structures;
    if ... there exists, as the minimum, a transition which may fire ... then
      begin
        Firings:= TRUE;
        while ...there exist transitions which are in alternative and/or in conflict ... do
          begin
            generate a pseudorandom index with uniform distribution for
            transitions which may fire in the frame of alternative and/or
            conflict net structures;
            execute the transition with respect to the generated index;
            delete all the transitions which are no more in alternative and/or
            in conflict from the proper list;
          end;
        execute all transitions which may have simple firings;
      end
    else
      Firings:= FALSE;
    until ...there is no transition which may fire ... ;
  end;
end;
```

For the sake of completeness, let's take a look at the procedure that executes a transition firing:

procedure ExecuteTrans(TransId);

{TransId is the identifier code of the transition to be executed}

begin

if ...transition has an associated algorithm ... **then** execute the algorithm;

decrease the marking of input places according to the firing rule and multiplicities of arcs;

increase the marking of output places according to the firing rule and multiplicities of arcs;

for all output places:

if ...place has an associated MO ... **then**

insert the associated MO in the list of active processes;

end;

References

1. G. Degli Antoni and G. Haus, "Music and Causality," *Proc. 1982 ICMC*, Computer Music Assoc., San Francisco, 1983, pp. 279–296.
2. S. Pope, "The Development of an Intelligent Composer's Assistant," *Proc. 1986 ICMC*, Computer Music Assoc., San Francisco, 1986.
3. C.A. Petri, "General Net Theory," *Proc. Joint IBM and Newcastle upon Tyne Seminar on Computer Systems Design*, 1976.
4. J.L. Peterson, *Petri Net Theory and the Modeling of Systems*, Prentice Hall, Englewood Cliffs, N.J., 1981.
5. R. Valk, "Self-Modifying Nets, a Natural Extension of Petri Nets," *ICALP 1978*, LNCS, No. 62, Springer-Verlag, Berlin, 1978, pp. 464–476.
6. V.E. Kotov, "An Algebra for Parallelism Based on Petri Nets," *MFCs 1978, Proc. 7th Symp.*, Springer-Verlag, Berlin, 1978.
7. G. Haus and A. Rodriguez, "Music Description and Processing by Petri Nets," *1988 Adv. Petri Nets*, LNCS, No. 340, Springer-Verlag, Berlin, 1989, pp. 175–199.
8. J.A. Goguen, "Complexity of Hierarchically Organized Systems and the Structure of Musical Experience," *UCLA Computer Science Dept. Quarterly*, Vol. 3, No. 4, 1975.
9. A. Bertoni, G. Haus, G. Mauri, and M. Torelli, "A Mathematical Model for Analyzing and Structuring Musical Texts," *Interface*, Vol. 7, No. 1, Swets & Zeitlinger B.V., Amsterdam, 1978, pp. 31–43.
10. D. Harel, A. Pnueli, J.P. Schmidt, and R. Sherman, "On the Formal Semantics of Statecharts," *Proc. Symp. Logic in Computer Science*, IEEE Computer Soc. Press, Los Alamitos, Calif., 1987, pp. 54–64.

Note: The musical and audio results produced by the system described in this article are an integral part of the project. For a reference of the publication in audio form and availability, and for a discussion of these musical results, see the *Discography*.

Goffredo Haus is scientific director of the Laboratorio di Informatica Musicale at the University of Milan. He does research in the university's Department of Computer Science and composes music by computer methods at both the score and sound level. He coordinates many research projects about scientific, technological, aesthetic, and multimedia topics cooperating with many Italian and international institutions and industries.

Haus is a founder of the Italian Association of Musical Informatics (AIMI) of which he has been a board member and president. He has also been board member of the Computer Music Association.

Alberto Sametti is the software manager of the Intelligent Music Workstation Project at the Laboratorio di Informatica Musicale at the University of Milan. He works on petri net applications to music and is a scientific consultant to Apple Computer on software development and validation.

