Procedural Generation of Several Instrument Music Pieces with Hierarchical Wave Function Collapse

Author: Raphael de Wolff, R.C.deWolff@student.tudelft.nl Supervisor and responsible professor: Rafael Bidarra

TU Delft

1. Background & motivation

Varga and Bidarra have developed a mixed-initiative procedural music generation model, which can generate music pieces with chords and a melody played by one instrument [1]. The model is based on a hierarchical version of WFC, a constraint-based procedural content generation algorithm.

Figure 1: Basic hierarchy for a music generator

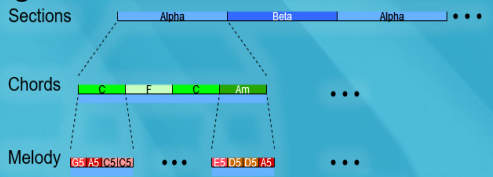

Exploring HWFC's capability of generating several instrument music pieces is a logical next step in WFC-based procedural music generation.

2. Research questions

How can an HWFC-based music generation model generate a piece of music that is played on multiple instruments in parallel?

1. What constraints can be introduced for interdependence between instruments?
2. In what ways can instruments playing chords and melodies be represented on a canvas?
3. How can the current HWFC algorithm be adapted for this several instrument canvas structure?
4. How good are the different canvas representations and algorithms in terms of efficiency and quality?

3. Method

Designing model: Extend Varga and Bidarra's model. Instruments play the same chord progression, so only the melody layer will change.

Designing constraints: Use Varga and Bidarra's hierarchical constraints. Also introduce intermelody constraints, which instruments impose on each other. Use the 4-part composition ruleset to design the intermelody constraints [2]

Evaluation procedure: We test different combinations of model variant, constraint sets and number of instruments on efficiency (average speed of a successful generation) and failure rate. We also listen to some successful generations and make comments on their quality and notable features.

4. Proposed model

Figure 2: Hierarchy example for a several instrument music generator

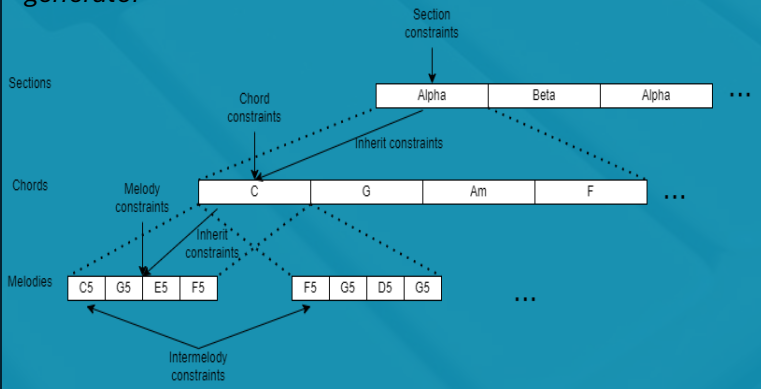

4 ways to collapse melody canvasses:

- **Naive**: Fully collapse canvasses sequentially.
- **Random**: Randomly choose a canvas to collapse one tile of.
- **Random k**: Randomly choose a canvas to collapse k tiles of.
- **Jam**: Collapse one canvas fully, then perform random (k) collapse on other canvasses.

5. Results

Efficiency: A single successful generation happens in the milliseconds to 10s of milliseconds, which for the purposes of the model is fast enough.

Failure rate: Failure rate heavily goes up when intermelody constraints are introduced. *Random collapse* is narrowly best at handling higher complexities, *Naive collapse* is clearly worst.

Quality: Successful generations sound somewhat like classical music. *Random k-* and *Jam collapse* sound most coherent, *Random collapse* least.

6. Conclusion

A HWFC-based music generation model can generate a (coherent) piece of music played on multiple instruments in parallel. It is unsure which model is best. Due to high failure rates among all models, all models would benefit from integrating backtracking.

7. Limitations

The evaluation of model quality is purely done by the researcher, and heavily subjective.

8. References

[1] Patrik Pál Varga and Rafael Bidarra. 2024. Harmony in Hierarchy: Mixed-Initiative Music Composition Inspired by WFC. Submitted for publication.

[2] Dave Smey. 2004. Important Rules for 4-Part Progressions. https://davesmey.com/theory/partwritingrules.pdf