Emulation of an Unconventional Model of Computation in Java

Aidan Delaney and Thomas J. Naughton

Department of Computer Science, National University of Ireland, Maynooth,
County Kildare, Ireland
{adelaney,tomn}@cs.may.ie

Abstract. This paper describes the emulation of an unconventional model of computation inspired by the field of optical computing. The model could be described as a random access machine with registers that hold continuous two-dimensional images. Our development employed a combination of eXtreme Programming, unit and integration testing with junit, and design patterns. In the final product we implemented a novel content-routing message passing system and have realised the first debugger for an optical computer programming language.

1 Introduction

Unconventional models of computation have been studied for many years. They differ to models such as the Turing machine [1], random-access machine [2], and the λ-calculus [3], usually in their atomic operations or data representation. Any model of computation that differs significantly to well-known universal models [1–3] could be regarded as unconventional. Examples of unconventional models of computation include models that compute over the reals [4], continuous-time models inspired by Newtonian mechanics [5], analog models inspired by neural networks [6], optical models inspired by Fourier optics [7, 8], models inspired by quantum mechanics [9], and biomolocular models inspired by molecular genetics [10, 11]. A snapshot of current research can be found in [12]. Studying unconventional models of computation allows us to re-evaluate our longstanding beliefs about computer science and gives us the opportunity to either falsify or strengthen the theses (Church-Turing thesis, Invariance thesis) on which computer science is built. Nature has often been the inspiration for unconventional models of computation [5–11]. Other models have been defined completely independently of natural processes [4], but the principal advantage of those inspired by Nature is their potential realisation. Computer scientists hope to one day build such devices, and harness the huge gains in computational complexity (and possibly computability) that they promise. This paper describes an emulator for an unconventional model inspired by optical physics. In this model all data is stored in infinite resolution images. The model has super-Turing capabilities and so cannot be fully emulated digitally [13]. Therefore, a programming methodology and programming tools were required that permitted necessary compromises to be made, as they presented themselves, during the development of an emulator.

We decided to use the eXtreme Programming (XP) paradigm. Design patterns were employed throughout. Java was chosen for our development as it offered modularity

and extensibility. Java complements agile development methodologies (such as XP) by being object-oriented and platform-independent, by having standard support libraries, and by taking care of memory management. Design pattern implementation is well documented in Java (the Java AWT itself implements several patterns). We were also able to take advantage of the XML parsers and DOM implementations available for Java. Unit testing in Java is straightforward using junit and its spin-offs such as XMLunit. We found that it was also possible to configure junit for integration testing. A long-term development goal of ours is to provide a web-based optical computer emulator, and Java's extensibility in this regard is exemplary.

In Sect. 2 we describe the model of computation. The emulator and novel content-routing message passing subsystem are explained in Sects. 3 and 4, respectively. The first (to our knowledge) debugger for an optical computing programming language is outlined in Sect. 5, and in Sect. 6 we describe unit and integration testing using junit.

2 Model of Computation

The continuous-space model of computation is inspired by the theory of Fourier optics. The model was developed for the analysis of (analog) Fourier optical computing architectures and algorithms, specifically pattern recognition and matrix algebra processors [14–16]. The functionality of the model is limited to operations routinely performed by optical scientists. The model operates in discrete timesteps over a finite number of two-dimensional (2D) images of finite size and infinite resolution. It can navigate, copy, and perform other optical operations on its images. A useful analogy would be to describe the model as a random access machine [2], without conditional branching and with registers that hold 2D functions. Universality, computational complexity gains, and super-Turing properties of the model have been demonstrated [13].

The memory of the machine consists of a finite 2D grid of images that hold both a program and the input. Each grid element holds a 2D infinite resolution complex-valued image of the form $f : \mathbb{R}_0^1 \times \mathbb{R}_0^1 \mapsto \mathbb{C}$ where $\mathbb{R}_0^1 = \{x : x \in \mathbb{R} \wedge 0 \leq x \leq 1\}$. There is a program start location **sta** and a small number of 'well-known' addresses labelled **a**, **b**, **c**, and so on. The two most basic operations available to the programmer, **ld** and **st** (parameterised by two column addresses and two row addresses), copy rectangular subsets of the grid into and out of image **a**, respectively. Upon such loading and storing the image information is rescaled to the full extent of the target location. The complete set of atomic operations is given in Fig. 1. Each instance of the machine is described by the quadruple $M = (D, L, I, P)$, where D is the grid dimensions, L is the locations (grid coordinates) of **sta** and the well-known addresses, I is the inputs and each of their locations, and P is the program instructions and each of their locations.

There are many ways to encode finite, countable, and uncountable sets as images. We have designed our number encoding scheme around an image that contains a high amplitude at its centre and zero everywhere else. Such an image encodes the symbol '1'. An empty image (zero amplitude everywhere) encodes '0'. Images can be combined using a stepwise rescaling technique (an image 'stack') or with a single rescale operation (an image 'list') to encode nonnegative integers in unary and binary notations. These concepts are illustrated in Fig. 2. The choice of encoding is usually driven

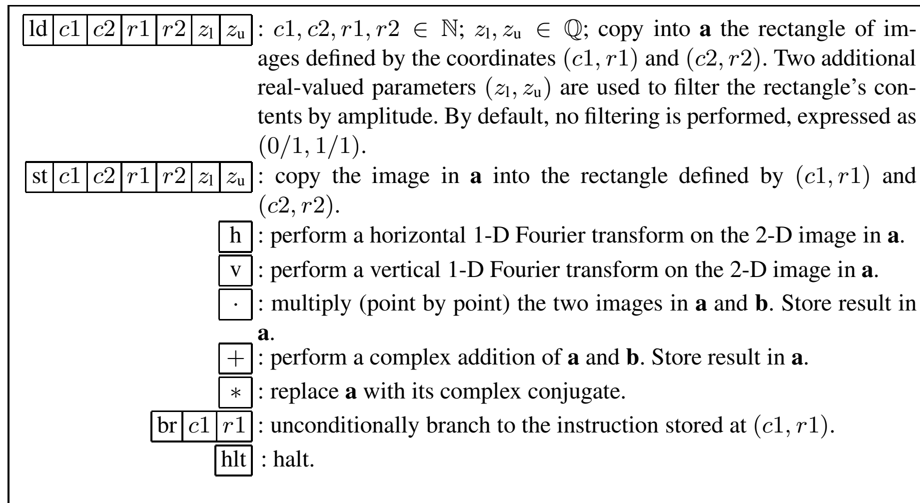

ld | $c1$ | $c2$ | $r1$ | $r2$ | z_l | z_u : $c1, c2, r1, r2 \in \mathbb{N}$; $z_l, z_u \in \mathbb{Q}$; copy into **a** the rectangle of images defined by the coordinates $(c1, r1)$ and $(c2, r2)$. Two additional real-valued parameters (z_l, z_u) are used to filter the rectangle's contents by amplitude. By default, no filtering is performed, expressed as $(0/1, 1/1)$.

st | $c1$ | $c2$ | $r1$ | $r2$ | z_l | z_u : copy the image in **a** into the rectangle defined by $(c1, r1)$ and $(c2, r2)$.

h : perform a horizontal 1-D Fourier transform on the 2-D image in **a**.

v : perform a vertical 1-D Fourier transform on the 2-D image in **a**.

$\cdot$: multiply (point by point) the two images in **a** and **b**. Store result in **a**.

$+$: perform a complex addition of **a** and **b**. Store result in **a**.

$*$: replace **a** with its complex conjugate.

br | $c1$ | $r1$: unconditionally branch to the instruction stored at $(c1, r1)$.

hlt : halt.

Fig. 1. A summary of the optical computing programming language [13]

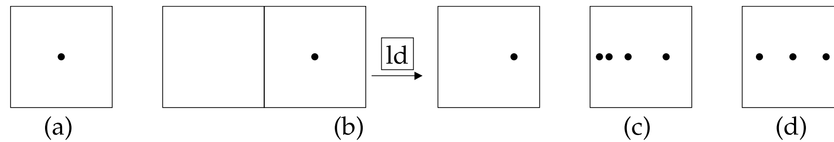

Fig. 2. Encoding numbers in images through the positioning of amplitude peaks. In the illustrations, the nonzero peaks are coloured black and the white areas denote zero amplitude. (a) The symbol '1' is encoded as a single peak in the centre of an image. (b) A 'stack' structure can be used to encode numbers in unary. We 'push' a symbol '1' onto an empty stack (the zero amplitude image) by loading both images simultaneously into **a** and rescaling to the dimensions of a single image. (c) The number 4 encoded in a stack. (d) The number 3 encoded in a 'list' structure

```
list   ::= l orientation body
stack ::= sh orientation body: | sv orientation body:
body ::= {complex_number}^positive_int | {complex_number}^*
```

Fig. 3. A fragment of the regular grammar describing a subset of the infinite-resolution images. The '^*' symbols code for an infinite sequence of peaks of amplitude specified by the preceding complex number

by computational complexity considerations. It is possible to enhance this encoding by allowing the peaks to have real-valued amplitudes, and by using both dimensions of the image.

3 Emulation of the Model

Emulation of the model of computation is not straightforward: the set of infinite-resolution images is uncountable. We must define an encoding scheme for a countably infinite subset of this set in order to emulate (at least partially) the model discretely. In addition, we must develop algorithms for operations over this subset.

The core of the emulator is the `Processor` subsystem, which handles the memory. Other subsystems are concerned with user input, message passing, processor control, and input verification. The user interface `OutputViewer` subsystem prints processor output and errors. It receives an XML DOM tree from `Processor` in the case of the former and an exception in the case of the latter. The other major subsystem, the `MessageCenter` communication subsystem, is explained in Sect. 4.

3.1 Memory

We decided to design an emulator that is independent of the image encoding. The core of the emulator, the `Processor` subsystem, executes all of the emulation operations. It stores the image grid in a 2D array, and uses an image iterator as an instruction pointer. Each image can itself contain a grid of images. This allows the emulator to load a rectangle of grid images into one image without explicitly needing to process the image data (thereby possibly introducing some loss of information due to rounding error). If an image is to be stored over a rectangle of grid images, it is copied to each image in turn and an appropriate viewbox is set in each image. The viewbox concept was borrowed from the scalable vector graphic specification. In practice, no scaling of images actually takes place in our emulator; images are resized by manipulating viewboxes.

3.2 Encoding of Images

The restricted set of images we decided to model is based on the encoding scheme shown in Fig. 2, which had been sufficient to prove some computability and computational complexity properties of the model [13]. A regular grammar (the core of which is shown in Fig. 3) was found to be sufficiently expressive. Each image contains either a word in the regular language or a finite-resolution pixelated bitmap (to also support all conventional raster images).

3.3 Extra Encoding Data

Extra encoding data was employed to emulate unary operations h, v, and $*$, and the binary operations $\cdot$ and $+$. It was used for the lazy evaluation of unary and binary operations over the encoded continuous images. For example, an image might be prepended with a symbol 'h' to indicate a horizontal Fourier transformation. This system is image representation-independent, be it a raster scheme, vector scheme, or (in our case) a regular grammar. Importantly, lazy evaluation of operations (possibly only when the final result is required by the user) means that defects due to rounding errors, etc., do not propagate through the computation and are only realised at a halt.

3.4 Cancellation Calculus

A cancellation calculus was developed to reduce redundant encoding data at each step in the computation, and to implement image operations. This calculus also has the ability to simplify concatenated viewbox representations. The calculus is given in the form of (i) a list of simplifying transformation rules, and (ii) algorithms for Fourier transformation and other image manipulation operations. As it is the only component of the emulator that actually physically realises operations, it is the only component specific to an image encoding scheme. Our emulator is therefore easily extensible. The range of infinite-resolution images can be increased by simply extending the grammar (to context-free, for example) and appropriately enhancing the rules and algorithms in the cancellation calculus.

4 Message Passing Subsystem

We describe the evolution, in terms of a sequence of refinements, of the `MessageCenter` from a Java `Observer-Observable` implementation to a generic content routing architecture with a simple design. We assume the precondition that an observable-observer relationship exists using the `Observer` and `Observable` classes provided by the Java language. The observer receives updates with differing content but with the same data structure. It differentiates updates based on content type rather than data type; by doing so we preclude the use of dynamic type checking to differentiate between updates. We implement an Adapter [17, p. 139] in the observable to wrap the update. A content-type (set by the observable) is provided as extra information by the adapter. Therefore, the observable can send all content packages to the observer, which accepts (or ignores) a packet based on its content type.

As we still use the Java `Observer-Observable` implementation, observable references will have been hard-coded into observers. In the next refinement, we replace hard-coded references to make the system more modular. The Mediator [17, p. 273] component is implemented as a well-known Singleton [17, p. 127]. This mediator forwards update requests from observables to observers. A reference to the mediator can be dynamically obtained by any component, as it is a well-known singleton.

A further refinement allows us to reduce the amount of code in content receivers. All observers receive all updates and have a long sequence of `if()..else()` statements

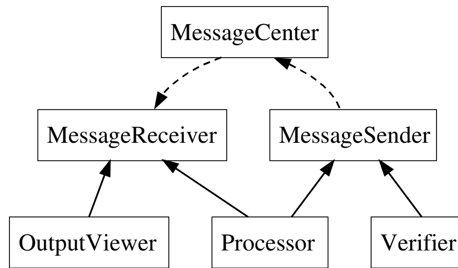

Fig. 4. A diagram of our generic message passing architecture

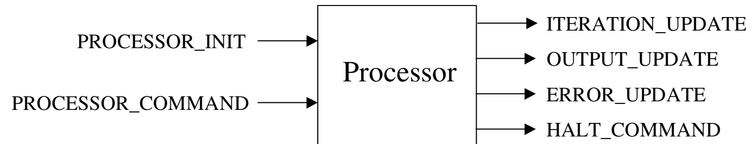

Fig. 5. A diagram of content-types passed into and out of the `Processor`

to filter out content types. Observers should elect to receive content they are interested in. The mediator accepts content sender registration requests from an observable and content receiver registration requests from observers. The mediator routes content from an observable to any observer interested in the content.

We also wanted to open our architecture for use as a generic message passing system. The architecture had to remain useful as an observable-observer update brokering mechanism. We created a `MessageSender` abstract class and `MessageReciever` interface. Any classes wishing to send a `Message` must extend `MessageSender` and register the content with the mediator. Any class wishing to receive a `Message` must implement `MessageReciever` and register to receive specific content. The mediator routes `Messages` from a `MessageSender` to possibly several `MessageRecievers` interested in the particular content. The final system is illustrated in Fig. 4.

5 Debugger

Our emulator has a processing core (`Processor`) and a user interface (`OutputViewer`) that communicate via our generic content routing system. The user is given the option of setting a breakpoint in the executing instructions and is informed when a breakpoint has been reached. In this context, our `Processor` and `OutputViewer` are the first two components of a Model-View-Control (MVC) pattern. We add a third component `IterationController` to act as a control. The `IterationController` registers to receive updates of the position of `Processor`'s instruction pointer. It also must send commands to the `Processor` instructing it to step and halt. `Processor`'s inputs and outputs are shown in Fig. 5. The timing chart for the MVC communication is shown in Fig. 6, and is explained next.

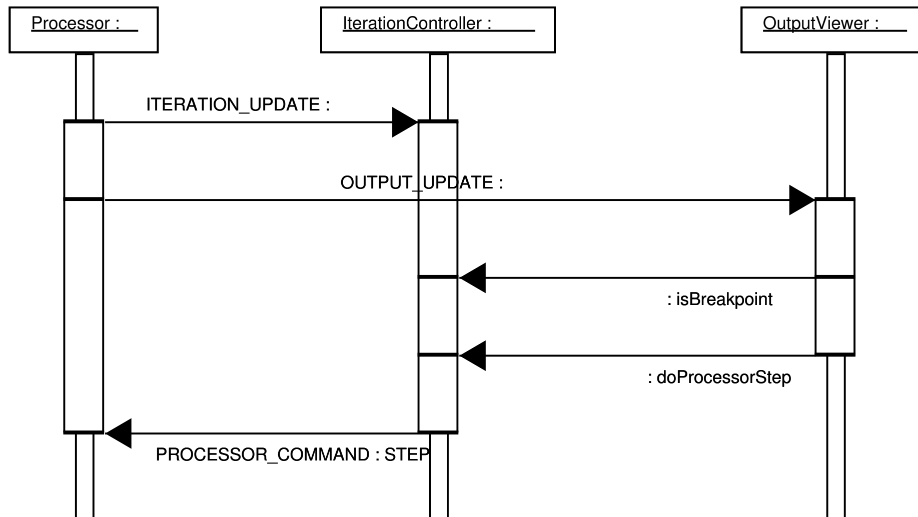

Fig. 6. A UML sequence diagram of our debugger. The role of the `MessageCenter` is hidden for clarity

The view (`OutputViewer`) will have allowed the user to insert a breakpoint (that is stored in the control). The control (`IterationController`) passes a Command [17, p. 233] to the model (`Processor`) instructing it to execute. After execution of each instruction the model updates the control with the current position of the instruction pointer. If the instruction pointer is at a breakpoint, the control provides this information to the view, otherwise the model is instructed to continue executing. If the view reads breakpoint information from the control, it displays a helpful message to the user, and waits for their response. On receipt of a response, the view instructs the control to command the model to continue execution.

At each breakpoint, the user has the option of stepping through the execution (repeating the sequence of operations in Fig. 6 for each instruction), or of continuous execution. If processing is paused and the view determines that it was not due to a breakpoint (or stepping) the content containing the current state of the image grid is regarded as the output of the computation.

6 Junit

We use junit to automatically run unit tests. Junit is a framework allowing programmers to implement the XP test-first practice with ease. Using junit normally involves subclassing `junit.framework.TestCase`. Numerous unit tests for a specific class may be written in this subclass (referred to as a test fixture).

For each of our classes we created a fixture. Our fixtures contain several tests, which are standard Java methods. The actual test result is compared with the expected result using the `assertTrue` and `assertEquals` methods provided by junit. A simple test on a class involves instantiating it and testing the validity of members available via

public `get()` and `set()` methods. We use the reflection "hack" provided by junit to allow easy extensibility of the test repository. The master test class contains only the name of the class containing test fixtures. Thus no changes are required to other test classes if a new test is added to a fixture.

6.1 Integration Testing with Junit

We identified a method for utilising junit for automating integration testing. The method was first applied to the message passing subsystem of `MessageSender`, `MessageCenter`, and `MessageReciever` (see Fig. 4), and secondly to `Processor`. These two subsystems were targeted as integration testing subjects because they (i) are two subsystems that are composed of several other classes, (ii) have a high dependency on classes within their subsystem (high cohesion), and (iii) have a low dependency on classes outside their subsystem (low coupling). By implementing internal classes in a junit fixture, we could subclass `MessageSender` and `MessageReciever`. The `MessageSender` could then pose as a `Verifier`, `Processor`, or `IterationController`. The `MessageReciever` could pose as a `Processor`, `OutputViewer`, or `IterationController`. In this fashion, message passing between modules could be tested. By implementing integration testing within the unit test framework, we could automate the whole verification testing process.

7 Conclusion

We have successfully emulated the partial functionality of a super-Turing unconventional model of computation. This required defining an encoding scheme for a subset of its infinite-resolution images. Java aided our agile development as it allowed us to group many classes with different functions together using an interface. This is seen in classes implementing the `MessageReciever` interface. It would not have made sense to use `MessageReciever` as an abstract class as receivers receive differing content types and must process these types differently. Design patterns, in addition to giving us trusted methods of solving common problems, allowed us to provide clear and concise explanations (for example, for the evolution of our message passing architecture).

Several XP practices were utilised. Practises such as collecting user stories and devising release and iteration plans were found to be advantageous. Building an automated unit test repository allowed us to take full advantage of the XP test-first practice, which we feel increased the pace of development. Our integration testing is effected using the same technique and so both unit and integration tests could become part of any future regression test suite.

Future work on this project will include enhancing the image encoding scheme. We have shown that only the cancellation calculus must be updated to increase the range of acceptable infinite-resolution images as our emulator is a general emulation framework. The software end-product is currently being used to design algorithms for implementation on Fourier optical computers.

Acknowledgments

The authors gratefully acknowledge the assistance of Damien Woods, and the Theoretical Aspects of Software Systems (TASS) research group, NUI Maynooth.

References

1. Turing, A.M.: On computable numbers, with an application to the Entscheidungsproblem. Proceedings of the London Mathematical Society ser. 2 **42** (1936) 230–265 Correction in vol. 43, pp. 544–546, 1937.
2. Shepherdson, J.C., Sturgis, H.E.: Computability of recursive functions. Journal of the Association for Computing Machinery **10** (1963) 217–255
3. Church, A.: An unsolvable problem of elementary number theory. American Journal of Mathematics **58** (1936) 345–363
4. Blum, L., Shub, M., Smale, S.: A theory of computation and complexity over the real numbers: NP-completeness, recursive functions and universal machines. Bulletin of the American Mathematical Society **21** (1989) 1–46
5. Moore, C.: Recursion theory on the reals and continuous-time computation. Theoretical Computer Science **162** (1996) 23–44
6. Siegelmann, H.T., Sontag, E.D.: Analog computation via neural networks. Theoretical Computer Science **131** (1994) 331–360
7. Naughton, T.J.: A model of computation for Fourier optical processors. In Lessard, R.A., Galstian, T., eds.: Optics in Computing 2000. Proceedings of SPIE vol. 4089, Quebec, Canada (2000) 24–34
8. Reif, J.H., Tyagi, A.: Efficient parallel algorithms for optical computing with the discrete Fourier transform (DFT) primitive. Applied Optics **36** (1997) 7327–7340
9. Deutsch, D.: Quantum theory, the Church-Turing principle, and the universal quantum computer. Proceedings of the Royal Society of London **A400** (1985) 97–117
10. Head, T.: Formal language theory and DNA: an analysis of the generative capacity of specific recombinant behaviors. Bull. Math. Biology **49** (1987) 737–759
11. Păun, G., Rozenberg, G., Salomaa, A.: DNA Computing : New Computing Paradigms. Springer-Verlag, Berlin (1998)
12. Antoniou, I., Calude, C.S., Dinneen, M.J., eds.: Unconventional Models of Computation – Proceedings of the Second International Conference on Unconventional Models of Computation (UMC'2K), Solvay Institutes, Brussels, Belgium, December 2000. Springer Series in Discrete Mathematics and Theoretical Computer Science, Springer, London (2001)
13. Naughton, T.J., Woods, D.: On the computational power of a continuous-space optical model of computation. In Margenstern, M., Rogozhin, Y., eds.: Third International Conference on Machines, Computations, and Universality (MCU 2001). Springer Lecture Notes in Computer Science vol. 2055, Chişinău, Moldova (2001) 288–299
14. Naughton, T., Javadpour, Z., Keating, J., Klíma, M., Rott, J.: General-purpose acousto-optic connectionist processor. Optical Engineering **38** (1999) 1170–1177
15. VanderLugt, A.: Signal detection by complex spatial filtering. IEEE Transactions on Information Theory **IT-10** (1964) 139–145
16. Weaver, C.S., Goodman, J.W.: A technique for optically convolving two functions. Applied Optics **5** (1966) 1248–1249
17. Gamma, E., Helm, R., Johnson, R., Vissides, J.: Design Patterns, Elements of Reusable Object-Oriented Software. Addison-Wesley, Upper Saddle River, NJ (1995)