



CRN++: Molecular Programming Language

Marko Vasic^(✉), David Soloveichik, and Sarfraz Khurshid

The University of Texas at Austin, Austin, USA
{vasic,david.soloveichik,khurshid}@utexas.edu

Abstract. Synthetic biology is a rapidly emerging research area, with expected wide-ranging impact in biology, nanofabrication, and medicine. A key technical challenge lies in embedding computation in molecular contexts where electronic micro-controllers cannot be inserted. This necessitates effective representation of computation using molecular components. While previous work established the Turing-completeness of chemical reactions, defining representations that are faithful, efficient, and practical remains challenging. This paper introduces *CRN++*, a new language for programming deterministic (mass-action) chemical kinetics to perform computation. We present its syntax and semantics, and build a compiler translating *CRN++* programs into chemical reactions, thereby laying the foundation of a comprehensive framework for molecular programming. Our language addresses the key challenge of embedding familiar imperative constructs into a set of chemical reactions happening simultaneously and manipulating real-valued concentrations. Although some deviation from ideal output value cannot be avoided, we develop methods to minimize the error, and implement error analysis tools. We demonstrate the feasibility of using *CRN++* on a suite of well-known algorithms for discrete and real-valued computation. *CRN++* can be easily extended to support new commands or chemical reaction implementations, and thus provides a foundation for developing more robust and practical molecular programs.

1 Introduction

A highly desired goal of synthetic biology is realizing a programmable chemical controller that can operate in molecular contexts incompatible with traditional electronics. In the same way that programming electronic computers is more convenient at a higher level of abstraction than that of individual flip-flops and logic circuits, we similarly expect molecular computation to admit specification via programming languages sufficiently abstracted from the hardware. This paper focuses on developing a compiler for a natural imperative programming language to a deterministic (mass-action) chemical reaction network implementing the desired algorithm. We do not directly make assumptions on how the resulting reactions would be implemented in chemistry. This could in principle be achieved by DNA strand displacement cascades [15], or other programmable chemical technologies such as the PEN toolbox [3].

Deterministic (mass-action) chemical kinetics is Turing universal [9], thus in principle allowing the implementation of arbitrary programs in chemistry. Turing universality was demonstrated by showing that arbitrary computation can be embedded in a class of polynomial ODEs [4], and then implementing these polynomial ODEs with mass-action chemical kinetics. While these results establish a sound theoretical foundation and show the power of chemistry for handling computation tasks in general, translating and performing specific computational tasks can lead to infeasibly large and complex sets of chemical reactions.

In this work we develop a programming paradigm for chemistry, based on the familiar imperative programming languages, with the aim of making molecular programming more intuitive, and efficient. Most commonly used programming languages such as C, Java and Python, are imperative in that they use statements that change a program’s state, with typical branching constructs such as if/else, loops, etc. Note that although CRNs are sometimes talked about as a programming language [7], they are difficult to program directly (it is even unfair to equate them with assembly language). In contrast, *CRN++* operates at a much higher level.

A mapping of imperative program logic to chemical reactions manipulating continuous concentrations poses various challenges that we must address. All reactions happen concurrently, making it difficult to represent sequential computation where, for example, the result of one operation is first computed and then used in another operation. Similarly, all branches of the program execution (i.e., if / else) are followed simultaneously to some degree.

We introduce the syntax and semantics of *CRN++*, which is, to our knowledge, the first imperative programming language which compiles to deterministic (mass-action) chemical reaction networks. *CRN++* has an extensible toolset including error analysis, as well as simulation framework [8]. We thus provide an automatic environment for simulating experiments based on *CRN++* programs.

A user specifies a *sequence* of statements, termed commands, to execute. Assignment, comparison, loops, conditional execution, and arithmetic operations are supported. The generated reactions are logically grouped into modules performing the corresponding command. Each module transforms initial species concentrations to their steady-state values which are the output of the module. We ensure that such modules are composable by preserving the input concentrations at the steady-state. Note that in mass-action chemistry all species occur with non-zero concentrations, and thus all reactions happen in parallel to some extent. To mimic sequential execution, we ensure that the reaction corresponding to the current command happens quickly, while other reactions are slow. For this we rely on a chemical oscillator in which the *clock* species oscillate between low and high concentrations, and sequential execution is achieved by catalyzing reactions with different clock species. To achieve conditional execution, we further need to ensure that the reactions corresponding to the correct execution branch readily occur, while those corresponding to other branches are inhibited. Our *cmp* module sets *flag* species to reflect the result of comparison, and these species catalyze the correct branch reactions.

Sequential execution as well as conditional branching leads to errors. Error is present because instructions (reactions) that should not execute, still do (at a smaller rate, of course). Moreover, the set of basic modules, such as addition, converge to the correct value only in the limit, thus computing approximately in finite time. To mitigate the error, we choose the set of modules to exhibit exponential (fast) convergence, and we provide a toolkit for error analysis and detection. Our tool quantifies the error, which can help a user identify the source of error, and guide the design of more accurate *CRN++* programs.

We demonstrate the expressiveness of our language by implementing and simulating common discrete algorithms such as greatest common divisor, integer division, finding integer square root, as well as real-valued (analog) algorithms such as computing Euler’s number and computing π , shown in the full version of this paper. We implement the *CRN++* compiler to reactions in Mathematica, and use the *CRNSimulator* package [8] to manipulate and simulate chemical reactions. *CRN++* is an extensible programming language allowing for easy addition of new modules; we release the open-source version [1] of the tool to enable others make use of it, and extend it further.

2 Examples

In this section we discuss the characteristics of chemical reaction networks (CRNs) through examples. First, the overall idea of computation in CRNs is presented, followed by example programs in *CRN++*. The focus is to give a high level idea of our technique, while later sections discuss internal details.

Although historically the focus of the study of CRNs was on understanding the behavior of naturally occurring biological reaction networks, recent advancements in DNA synthesis coupled with general methods for realizing arbitrary CRNs with DNA strand displacement cascades [15] opened the path to engineering with chemical reactions. In this work we are not interested in a way to engineer the molecules implementing a reaction but focus on reaction behavior and dynamics. We abstract away molecule implementation information and denote molecular species with letters (e.g. *A*).

Molecular systems exhibit complex behaviors governed by chemical reactions. To give a formal notation of chemical reaction networks, consider the CRN 1:

CRN 1 Example chemical reaction network



The CRN 1 consists of two reactions. A chemical reaction is defined with *reactants* (left side), *products* (right side), and *rate constant* which quantifies the rate at which reactants interact to produce products. To illustrate this, reaction 1 is composed of *reactants* = $\{A, B\}$, *products* = $\{A, B, C\}$, and rate constant

$k = 1$. Since most reactions in *CRN++* have the rate constant equal to 1, from now on we drop the rate constant when writing reactions, unless it is different than 1. Note that multiple molecules of same species can be in a list of reactants (analogously for products); to support this we use the multiset notation. As an example, to describe reaction: $A + A \rightarrow B$ we write $reactants = \{A^2\}$, where the upper index (2) represents multiplicity (number of occurrences).

It may seem that a molecule of C is produced out of nothing in reaction 1, since the multiset of reactants is a submultiset of the products. This represents a level of abstraction where *fuel* species that drive the reaction are abstracted away (i.e., the first reaction corresponds to $F + A + B \rightarrow A + B + C$). Making this assumption allows us to focus on the computationally relevant species. The choice to use general (non-mass/energy preserving) CRNs is an established convention for DNA strand displacement cascades [15].

When the molecular counts of all species are large, and the solution is “well-mixed”, the dynamics of the system can be described by ordinary differential equations (mass-action kinetics). Molecular concentrations are quantified by a system of ODEs, where concentration of each species is characterized by an ODE:

$$\frac{d[S]}{dt} = \sum_{\forall rxn \in CRN} k(rxn) \cdot netChange(S, rxn) \cdot \prod_{\forall R \in reactants(rxn)} [R]^{m_{rxn}(R)}(t)$$

This ODE characterizes concentration of species S ($[S]$), in a given CRN. The right side is a sum over reactions in the CRN, where $k(rxn)$ is a rate of reaction rxn , and $netChange(S, rxn)$ is a net change of molecules of S upon triggering of rxn (can be negative). Concentration of a reactant R in time is written $[R](t)$, while $m_{rxn}(R)$ is the multiplicity of reactant R in reaction rxn . To illustrate the general formula, the set of ODEs characterizing CRN 1 is:

$$\frac{d[A]}{dt} = 0, \frac{d[B]}{dt} = 0, \frac{d[C]}{dt} = [A](t) \cdot [B](t) - [C](t)$$

The $[A]$ and $[B]$ are constant (derivatives zero); thus $\frac{d[C]}{dt} = [A](0) \cdot [B](0) - [C](t)$. From this equality follows that $[C](t)$ is increasing when smaller than $[A](0) \cdot [B](0)$, decreasing in the opposite case, and does not change when $[C](t) = [A](0) \cdot [B](0)$. Thus this system has a global stable steady-state $[C] = [A](0) \cdot [B](0)$. We say that this module computes multiplication, due to the relation between initial concentrations and concentrations at the steady state.

We simulate and plot the dynamics of the multiplication CRN, as shown in Fig. 1. Initial concentrations of A and B are 6 and 2, respectively, while the concentration of C approaches value 12. Note that the exact value defined by the steady state ($[C](t) = 12$) is reached only at the limit of time going to infinity. Since the computation has to be done in finite time, the presence of error is unavoidable. This

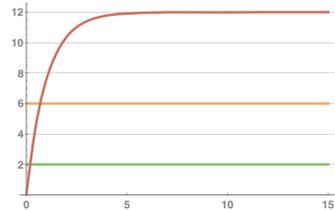


Fig. 1. Multiplication CRN. $[A]$ shown in orange, $[B]$ in green, and $[C]$ in red. (Color figure online)

error raises challenging issues with programming in chemistry, and necessitates techniques for controlling it. One crucial property that determines the error is the convergence speed of the module. The multiplication command in *CRN++* is implemented through the above module, following the design principles of *convergence speed* and *composability* described in Sect. 3. Chemical reactions are abstracted away from a user who can simply write $mul[a, b, c]$ to multiply.

CRN++ is an imperative language, and as such supports sequential execution. Note that even a simple operation of multiplying and storing into the same variable, e.g. $A := A * B$, requires support for sequential execution. We use operator “:=” to relate input and output concentrations; $A := A * B$ denotes that $[A](t)$ converges to $[A](0) * [B](0)$. The above implementation of the *mul* module necessarily assumes that the output species is different from the input species. Otherwise, $mul[a, b, a]$ goes to infinity or 0 depending on the value of B . To implement $A := A * B$, we split the computation into two sequential steps: (1) $C := A * B$, (2) $A := C$. To multiply we use the *mul* module described above. For the assignment we use the load module (*ld*). To ensure the assignment executes after the multiplication, we catalyze the two modules with the clock species that reach their high values in different phases of the oscillator. Importantly, the chemical oscillator and clock species are abstracted away from a user, who simply uses the *step* construct to order reactions: $step[\{mul[a, b, c]\}, step[\{ld[c, a]\}]$.

One of the basic blocks of programming languages are conditional branches, executing upon success of a precondition. Similarly to implementing sequential operations, we implement conditional execution by activating (through catalysis) some reactions and deactivating others, depending on a result of condition. Since no species can be driven to 0 in finite time¹, all branches of condition will be active to some extent, which makes this an interesting source of errors without direct analogy in digital electronics. In contrast to sequential computation catalyzed by clock species, conditional blocks are catalyzed by *flag* species. The flag species have high and low values that reflect the result of the comparison. Our *cmp* module sets the flag species to reflect the result of the comparison. In the following example we demonstrate the usage of *cmp* module and conditional execution.

```

1: procedure GCD( $a, b$ )
2:   while  $a \neq b$  do
3:     if  $a > b$  then
4:        $a \leftarrow a - b$ 
5:     else
6:        $b \leftarrow b - a$ 
7:     end if
8:   end while
9:   return  $a$ 
10: end procedure

```

Fig. 2. Euclid’s algorithm for computing GCD.

¹ Although certain pathological CRNs can drive concentrations to infinity in finite time (e.g., $2A \rightarrow 3A$), and thereby drive certain other species to 0 in finite time (e.g., with an additional $B + A \rightarrow A$), these cases cannot be implemented with any reasonable chemistry.

To demonstrate the expressiveness of our language we showcase the implementation of Euclid’s algorithm (Fig. 2) to compute the greatest common divisor (GCD) of a two numbers. The GCD is computed by subtracting the smaller of the values from the larger one until they become equal.

Figure 3a shows the implementation of Euclid’s algorithm in *CRN++*. Lines 2–3 define the initial concentrations of species a , b ; $a0$ and $b0$ represent the values for which GCD is computed. To order the execution, the *step* construct is used. Multiple instructions that do not conflict with each other can be part of the same step and they are executed in parallel. In the first step a and b are stored into temporary variables and compared, setting the flag species to reflect the result of the comparison. The second step uses the result of the previous comparison, and effectively stores $a - b$ into a if $a > b$, and vice versa. Since the same species cannot be used as both input and output to *sub* module, temporary variables are used (*atmp* and *btmp*). Steps repeatedly execute due to the oscillatory behavior of the clock species, thus implementing looping behavior by default; the steps can be viewed as being inside of the ‘forever’ loop. *CRN++*, in addition to the language and compiler to chemical reactions, is connected to the simulation backend that enables convenient testing for correctness. We show simulation of the GCD program in Fig. 3b where $\text{GCD}(32,12)$ is computed. Steps repeatedly trigger causing a and b to converge to the correct result after a couple iterations.

```

1 crn = {
2   conc[a,a0],
3   conc[b,b0],
4   step[{
5     ld[a, atmp],
6     ld[b, btmp],
7     cmp[a,b]
8   }],
9   step[{
10    ifGT[{ sub[atmp,btmp,a] }],
11    ifLT[{ sub[btmp,atmp,b] }]}
12 ]}
13 };

```

(a) GCD implementation

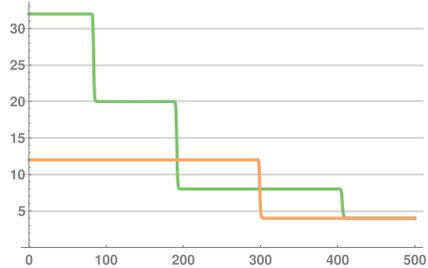
(b) Dynamic simulation of the GCD program for $a_0 = 32$, $b_0 = 12$. Concentrations of a (green), and b (orange) are shown in function of time.

Fig. 3. Implementation of Euclid’s algorithm for computing GCD in *CRN++* (left), simulation results of the implementation (right).

In addition, we implement a set of algorithms in (a) discrete space—counter, integer division, integer square root, as well as in (b) continuous space, by implementing *CRN++* programs that approximate value of *Euler’s* constant and π . These examples are shown in the full version of this paper.

3 Technique

This section explains *CRN++*, both the underlying constructs used to build it, as well as high level primitives that represent the language itself. We start by presenting high-level modules that are at the core of *CRN++* (Sect. 3.1), followed by explanation of how the sequential behavior is achieved (Sect. 3.2), after which we give an overview of *CRN++* grammar (Sect. 3.3), and finally we discuss the error detection and analysis tools we provide (Sect. 3.4).

3.1 Modules

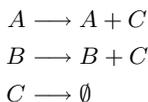
Modules represent the core of *CRN++*, and in their form are somewhat analogous to the instruction set architecture (ISA) in machine languages. Modules implement basic operations such as load, add, subtract, multiply, compare. We provide the exhaustive list of modules in Table 1 in Appendix A. Importantly, *CRN++* is extensible, and supports easy addition of new modules.

There are multiple ways of computing addition and other operations in chemistry. As mentioned in the previous section, our implementation choice is led by two basic principles: (a) convergence speed, (b) composability.

3.1.1 Convergence Speed

Consider CRN implementing addition:

CRN 2 Addition CRN (inputs preserved). Inputs: A and B , output: C .



By solving the system of ODEs that characterize the concentration of C we get the following equation: $[C](t) = [A] + [B] + ([C](0) - [A] - [B]) \cdot e^{-t}$.

$[C](t)$ is concentration of species C at time t ; accordingly $[C](0)$ is initial concentration. Since $[A](t)$ and $[B](t)$ are constant we simply write $[A]$ and $[B]$. From the equation above it follows that $[C]$ converges to the value $[A] + [B]$, and thus we say the CRN performs addition. To consider the convergence speed we look at the non-constant part of the equation. Due to the factor e^{-t} the decrease of the non-constant part is exponential, thus we say that the CRN exhibits *exponential* convergence speed. The convergence speed is of great importance, since it directly affects computation error; the sooner reaction converges the sooner it approaches the correct value.

3.1.2 Composability There are alternative ways to implement addition and have exponential convergence speed:

CRN 3 Addition CRN (destructs inputs). Inputs: A and B , output: C .

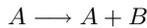


For easier discussion, let us call the initial addition module $CRNAdd_1$ (CRN 2), and the one above $CRNAdd_2$ (CRN 3). To compute $E := (A * B) + D$ we combine the mul module (CRN 1), computing $C := A * B$, with an addition module, computing $E := C + D$. If $CRNAdd_1$ is used, multiplication converges to the correct value, after which $CRNAdd_1$ has correct values at its inputs and converges to the correct value of E . Before the multiplication converges and C becomes equal to $A * B$, reactions of $CRNAdd_1$ trigger, but since the module is input-preserving they do not affect steady state of the multiplication module. However, $CRNAdd_2$ consumes its inputs, and the composition will give an incorrect result. The mul CRN constantly drives C to value $A * B$, and will keep refilling inputs to the $CRNAdd_2$, causing the wrong result. This is the reason $CRNAdd_1$ is preferred over $CRNAdd_2$. Moreover, the composed CRN of the mul module and $CRNAdd_1$ exhibits the exponential convergence speed, and has a unique stable steady state, the formal proof can be found in work by Buisman et al. [5].

We have set up the two main design criteria (convergence speed and composability) for the modules, and we next describe the core modules of $CRN++$.

3.1.3 Ld Module Loads the value from source (first argument) into a destination (second argument). The CRN used for load operation is following:

CRN 4 Load CRN

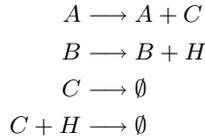


A is the input and B is the output species. This module, similar to add , has exponential convergence speed [5]. In addition, the concentration of input species is constant, thus ensuring composability.

3.1.4 Add Module Adds two values (first and second argument) and stores the result into destination (third argument). The Add CRN is shown in CRN 2; its convergence speed and composability are already discussed.

3.1.5 Sub Module Subtracts the second input value from the first and stores into the destination (third argument).

CRN 5 Subtraction CRN



The above CRN was generated via evolutionary algorithms [5]; by analyzing its system of ODEs, the network computes subtraction. Input species A and B are not affected and the property of composability is satisfied. Neither we nor Buisman et al. found the analytical solution; however, our simulation results show that the module converges exponentially quickly unless $A = B$ (see the Alternative Design subsection of the *cmp* module below for an analogous, easy to analyze case). In a case inputs, A and B , are close to each other the computation error is higher. The error evaluation tools (Sect. 3.4) help in detecting and analyzing problematic cases (e.g., where A and B are close), thus enabling a user to redesign the CRN. In our examples, A and B usually differ by at least 1. Runtime assertions in the simulation package that automatically notify a user about these kind of problems would help identify the source of the error. Note that many algorithms can be refactored to reduce the error (see Sect. 5).

3.1.6 Mul Module Multiplies inputs (first and second argument) and stores into destination (third argument). The multiplication CRN is shown in Sect. 2. This CRN does not affect inputs and has exponential convergence speed [5].

We have presented modules for performing arithmetic operations (*ld*, *add*, *sub*, *mul*). These modules are implemented within a single *step*. Multiple modules can be executed in parallel within a single *step* as long as there is no cyclic dependence between species: for example *mul*[a,b,c] and *add*[c,d,a] forms a cycle, the output of the *mul* is input to the *add*, and vice versa. Also, the CRN implementation imposes the restriction that same species cannot be used as both input and output to the same module. We now introduce the *cmp* module providing for conditional execution, which is executed in two *steps*.

3.1.7 Cmp Module Compares the two values, and produces signals (flag species) informing which value is greater or if they are equal.

Alternative Designs. Before explaining our implementation of comparison we discuss alternative implementations, and point out design decisions that lead to the current implementation. One of more obvious ways to implement comparison is using the reaction: $A + B \rightarrow \emptyset$.

If initially $[A] > [B]$, than at the equilibrium all molecules of B interact with A leaving only A ; case $[B] > [A]$ is analogous. The proposed reaction can be used for conditional execution, by using A and B catalytically in reactions that should execute when $[A] > [B]$ and $[B] > [A]$, respectively. The comparison reaction should execute before conditional execution is plausible, thus the comparison is done in a step before the conditionally executed reactions.

The comparison module proposed above does not preserve inputs, and thus it is not composable. This imposes the restriction that in the step in which comparison is used no other module uses the compared values. Our *cmp* module does not have this restriction.

We analyze the ODE describing this CRN to evaluate the convergence speed. Since the amount of B decreases with the same speed as A , we can express $[B](t) = [A](t) + D_0$, where $D_0 = [B](0) - [A](0)$. The following holds:

$$\frac{d[A]}{dt} = -[A](t) * ([A](t) + D_0) \implies [A](t) = \frac{[A](0)D_0}{-[A](0) + [A](0)e^{D_0t} + D_0e^{D_0t}}$$

If $D_0 > 0$ ($[B](0) > [A](0)$) terms with exponential factors tend to infinity, and $[A]$ to zero. Conversely, when $D_0 < 0$, exponential factors converge to zero, and $[A]$ to $-D_0$. A converges exponentially, unless A and B are equal at the beginning ($D_0 = 0$); then the dynamics of A is described with: $[A](t) = \frac{[A](0)}{1+[A](0)t}$.

In conclusion, the module converges fast (exponential speed) when operands are different, while the module converges slow (linear speed) when operands are equal (or close to each other). The linear convergence speed is yet another problem that lead to sub-optimal performance of this module. Recall that the comparison module drives the flag species which then catalyze branches that should execute, thus having a chained effect. It is of great importance to have a reliable comparison module.

Lastly, to detect equality with the above proposed module, absence of a species needs to be detected, since both values are driven to zero in a case of equality. Detecting the absence of species in chemistry is itself non-trivial and error-prone. There are several approaches based on so-called *absence indicators*. Generally speaking, the absence indicator for A is produced at a constant rate and gets degraded by A . The absence indicator has to be produced slowly, or else it will be present in non-negligible concentration even if A is present. The absence indicators in the literature rely on a difference between rate constants of several orders of magnitude. The relatively slow dynamics of the production of the absence indicator lead to a fair amount of error affecting the computation, and necessitate slowing down the clock (i.e., the whole computation) to work properly.

Our Design. The *cmp* module is implemented using two sequentially executed sets of reactions, which trigger in consecutive clock phases. In the first phase, the inputs (X and Y) are mapped to flag species $XGTY$ and $XLTY$. Values are mapped to the range $[0-1]$, by setting the initial concentrations of $XGTY$ and $XLTY$ to 0.5. If, for example, $[X] = 80$ and $[Y] = 20$, signal species $XGTY$ and $XLTY$ converge to 0.8 and 0.2, respectively. The mapping is done in order

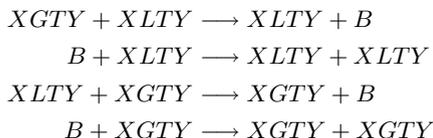
to preserve original values of the inputs, since the next phase of comparison consumes the compared values (flags), thus mapping allows the inputs to be used freely in other instructions. The mapping CRN is shown in CRN 6, and exhibits exponential convergence speed according to our analysis.

CRN 6 CRN for mapping compared values



The goal of the second phase of comparison is to detect which value is greater. We use a chemical *Approximate Majority* (AM) algorithm [6] to detect if $XGTY$ or $XLTY$ is in the majority. All molecules of the less populous species convert to the more populous species. AM reactions are:

CRN 7 Approximate Majority CRN



The majority algorithm causes convergence of $XGTY$ to 1 and $XLTY$ to 0 when $X > Y$, and vice versa. The species $XGTY$ are used as a catalysts in reactions that execute when $X > Y$, and the species $XLTY$ for the opposite case. The AM network has been studied in the stochastic context (stochastic CRNs) and is known to converge quickly, even when inputs are close [2].

Equality Checking. Due to the always present error in chemical computation, checking for equality is actually approximate-equality checking. Consider having a chemical program with real values, then if the values are close to each other it is impossible to tell if they are actually equal but affected with error, or they represent different real valued signals. Due to this issue, while comparing for equality is impossible, we compare for ϵ -range equality, where ϵ can be arbitrarily small. For discrete algorithms we use equality checking with $\epsilon = 0.5$, allowing easy comparison of the integer values (e.g., values in range (2.5, 3.5) are considered to be equal to 2). To support equality checking we compare $x + \epsilon$ with y (generating signals $XGTY$ and $XLTY$), and at the same time compare $y + \epsilon$ with x (generating signals $YGTX$ and $YLTX$). Combining the signals of the two comparisons gives the desired result: If $X = Y$, signal $XGTY$ is high ($XLTY$ low) and $YGTX$ is high ($YLTX$ low) due to the added offset. To execute a reaction upon equality both $XGTY$ and $YGTX$ are used catalytically. If $X > Y$, signal $XGTY$ is high ($XLTY$ low) and $YLTX$ is high ($YGTX$ low), so both $XGTY$ and $YLTX$ should be used catalytically. Symmetrically for $X < Y$,

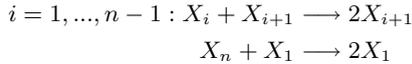
both $XLTY$ and $YGTX$ are used catalytically. Note that unlike in the previously proposed comparison module, this module does not use absence checks and absence indicators, and as such is more reliable in time-constrained environment. After calling cmp in a step, programmer can use $ifGT$, $ifGE$, $ifEQ$, $ifLT$, $ifLE$ in subsequent steps to conditionally execute reactions. Note that the flags are active until the next call to the cmp module.

3.2 Sequential Execution

$CRN++$ allows programming in a sequential manner, despite the intrinsically parallel nature of CRNs. To model sequential execution in CRNs there is a need to isolate two reactions from co-occurring, and control the order in which they happen. The key construct we rely on to achieve these goals is a chemical oscillator.

A chemical oscillator is a CRN in which the concentrations of species oscillate between low and high values. The oscillatory CRN [12] we use is described with a following set of reactions:

CRN 8 Oscillator CRN



Concentrations of the clock species (X_i) oscillate (see Fig. 4). Different clock species have different oscillation phase and reach minimum and maximum at different times. To control the rate at which a reaction fires, clock species are added as both reactant and product (catalyst), in that way preventing reactions from co-occurring and ordering them (see CRN 9). While overlap between the clock species exists, it is small and thus enables sequential execution. To ensure the small overlap, in $CRN++$ we use every third clock species, i.e. X_3 , X_6 , X_9 , etc., to catalyze the reactions that execute at different time moments.

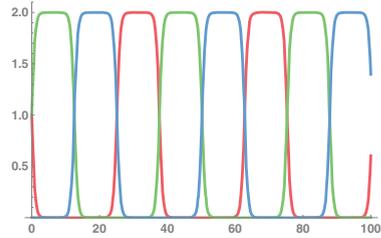


Fig. 4. Chemical oscillator containing 3 species: X_1 (red), X_2 (green), and X_3 (blue). (Color figure online)

CRN 9 Ordering reactions: original reactions (left), ordered (right).



The chemical oscillator is abstracted from a *CRN++* user, who can order reactions using the *step* construct. Reactions in different steps are isolated from each other through clock species acting catalytically.

Non-conflicting instructions can be part of the same step. Splitting instructions across multiple steps is needed in a case of (a) conditional execution—comparison needs to be done before conditional execution is possible; (b) reading and writing to a same species—this is not possible within the step (as discussed earlier), and requires temporal ordering. The number of clock species used is determined by the number of **step** instructions. Each step requires three clock species, with the exception of steps in which *cmp* module is used, for which six clock species are used. The oscillatory behavior of the clock species causes steps to get repeated eventually, causing the loop-like behavior.

3.3 Grammar

Listing 1.1 is an overview of the *CRN++* grammar. At its root the CRN contains a list of *RootSs*, where *RootS* can be either *ConcS* – defines initial concentration of species, *RxnS* – defines a reaction, *ArithmeticS* – performs arithmetic operation, and *StepS* – orders execution. Furthermore, *StepS* is divided into a list of *NestedSs*, where each *NestedS* is either *RxnS*, *ArithmeticS*, *CmpS* – performs comparison, or *ConditionalS*. *ConditionalS* conditionally executes a block based on the result of a previous comparison. Note that the comparison should be done in a step prior to the conditional execution. Based on the result of the comparison, if the first operand is *greater than*, *greater or equal*, *equal*, *less or equal*, *less than* the second operand, conditional block *ifGT*, *ifGE*, *ifEQ*, *ifLT*, *ifLE* is executed.

The grammar can be easily extended; e.g., new arithmetic modules can be added to the list of *ArithmeticS* nonterminals. Also, the *CRN++* grammar allows for easy addition of *ifAbsent* conditional statements that can be used to implement the asynchronous programs, in that way enabling comparison between the asynchronous and synchronous programming paradigms.

3.4 Error Evaluation

Programming chemistry is inherently error-prone. We identify three specific sources of error in *CRN++*. First, CRNs converge asymptotically—only in the limit is the correct value reached—thus leaving certain amount of error in a finite time. Second, we cannot completely turn off modules which are not supposed to be currently executing, whether they belong to another sequential step, or to another branch of execution. In addition, comparison has to take into account possible error in the compared values.

Our design decisions were based on minimizing the error; however since error cannot be avoided altogether, we provide a toolkit that helps in error analysis and guiding the CRN (program) design. Using the tool, users can, for any species of interest, track the difference between the correct value, and the (simulated) value in chemistry. For example, if operation $add[a,b,c]$ is executed in a step, than $c = a + b$ is expected in the following step. *CRN++* allows measuring the difference between the expected $c = a + b$, and actual simulation value. This helps users analyze the error, and detect if the error builds up over time.

We analyze the value of operand a from GCD example Fig. 3, and plot the error in Fig. 5. In Fig. 5, the x-axis represents time, while the y-axis shows the difference between expected and actual value of a . Note that the error is sufficiently small that the algorithm executes correctly throughout the analyzed time. The error is not constant, which opens interesting questions of correlating the error with instructions in the program. To correlate error with program instructions we examine the GCD simulation (Fig. 3b). By looking at the time axis, it is easy to connect the first two spikes of the error with the subtraction of a .

We provide the error evaluation framework with the vision of it being a guiding element for programming in *CRN++*. We found this technique particularly useful for validation of programs, analyzing the error, understanding the sources of error, and redesigning the CRN for correctness.

4 Related Work

Computational Power of Chemical Reaction Networks. Previous research demonstrated techniques of achieving complex behaviors in chemistry, such as: computing algebraic functions [5], polynomials [14], implementing logic gates [13]. Moreover, the Turing completeness of chemistry has been proven,

```

<Crn> ::= 'crn= {' <RootSList> '}'
<RootSList> ::= <RootS>
              | <RootS> ',' <RootSList>
<RootS> ::= <ConcS>
           | <RxnS>
           | <ArithmeticS>
           | <StepS>
<ConcS> ::= 'conc[' <species> ', ' <number> ':' ]'
<RxnS> ::=
          'rxn[' <Expr> ', ' <Expr> ', ' <number> ':' ]'
<ArithmeticS> ::=
          'ld [' <species> ', ' <species> ]
          | 'add [' <species> ', ' <species> ', ' <species> ]
          | 'sub [' <species> ', ' <species> ', ' <species> ]
          | 'mul [' <species> ', ' <species> ', ' <species> ]
<CmpS> ::= 'cmp [' <species> ', ' <species> ]
<StepS> ::= 'step [' <NestedSList> ']'
<NestedSList> ::= <NestedS>
                 | <NestedS> ',' <NestedSList>
<NestedS> ::= <RxnS>
             | <ArithmeticS>
             | <CmpS>
             | <ConditionalS>
<ConditionalS> ::= 'ifGT [' <NestedSList> ':' ]'
                  | 'ifGE [' <NestedSList> ':' ]'
                  | 'ifEQ [' <NestedSList> ':' ]'
                  | 'ifLT [' <NestedSList> ':' ]'
                  | 'ifLE [' <NestedSList> ':' ]'
<Expr> ::= <species> { '+' <species> }

```

Listing 1.1: *CRN++* Grammar

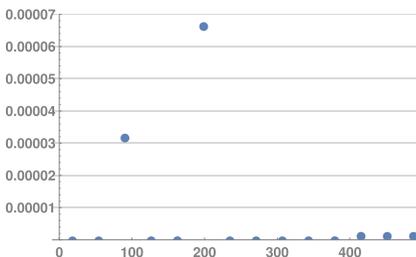


Fig. 5. Error evaluation of species a from GCD program.

using the strategy of implementing polynomial ODEs (which have been previously shown to be Turing universal) in mass-action chemical kinetics [9]. Even though Turing complete, this translation to chemistry can result in infeasibly complex chemical reaction networks, which motivates other, more direct methods.

Modular Reactions. Adding even a single reaction to a CRN can completely change its dynamics, which makes the design process challenging. The idea of ‘composable’ reactions seeks a set of reactions that can be composed in a well-defined manner to implement more complex behaviors. Buisman et al. [5] compute algebraic expressions by designing the core modules that implement basic arithmetic operations, which can be further composed to achieve more complex tasks. Our goal is to make modular designs, and we follow some of the proposed design principles for achieving the goal, such as input-preserving CRNs.

Synchronous Computation. Previous work utilized synchronous logic to achieve complex tasks. Soloveichik et al. implement state machines in chemistry by relying on clock species [15]. We use the same technique, where we add clock species acting catalytically to order reactions. Jiang et al. [11], also relying on clock species, design a model of memory in chemistry to support sequential computation, demonstrating their technique on examples of a binary counter and a fast Fourier transform (FFT). Previous work shows the promise of programming synchronous logic in reactions, which we advance by providing an explicit programming language and framework for designing and testing wide-range of programs. Maybe we should criticize previous work. Extend the paragraph by adding more comparison.

Asynchronous Computation. Recall, an absence indicator is a species that is present in high concentration when a target species is present in low concentration. Absence indicators can be used to drive a reaction when a particular reaction has finished, providing a method for executing modules in desired order. Huang et al. [10] use absence indicators to implement algorithms such as integer division and GCD. Their method requires two reaction rates, ‘fast’ and ‘slow’, where the fast rate needs to be orders (2–3) of magnitude larger to ensure the proper function of the system. Since, in practice, biochemical systems allow for a restricted range of reaction rates, requiring a large spectrum of rates slows down the computation when the computation speed is dictated by the slow rates. In contrast, we allow all reactions to take the same (or comparable) rate constants. While the goal of our work is not to compare asynchronous and synchronous computation, we mention insights and intuition of their differences, which we gained through empirical studies. First, absence indicators are not robust, and typically require fine tuning to get the system right. Second, error detection is easier with synchronous logic. Since all operations follow the clock signal, there is a direct mapping from a time moment to a command that is executing, which provides a way to check correctness of a system at any point of time. Finally, we provide a framework for implementing molecular programs which is easily extensible, and can be used to compare synchronous and asynchronous logic.

We include support for absence indicators through the *ifAbsent* construct, thus allowing easy comparison of the two paradigms.

5 Discussion and Conclusions

There are multiple ways in which we can further improve *CRN++*. Note that currently every high-level module is mapped to exactly one CRN implementing the operation. Letting the tool decide which implementation to use in different contexts could boost the performance. For example, the described modules have a useful property of preserving inputs, but that property might not be needed in every case. If the input preserving property is redundant, *CRN++* could choose to use the more optimized version (for example the more compact subtraction CRN discussed above). Also, we could provide a more flexible programming experience by (a) letting the compiler automatically schedule instructions to different steps (instead of the explicit `step` construct); (b) allowing the same species as both input and output of a module and automatically generate the additional instructions.

We plan to further explore the support for nested loops in *CRN++*. Currently nested loops can be mimicked through conditional execution: the loop condition is computed through comparison and the main loop conditionally executes the instructions of the desired loop. Besides explicit support for nested loops, future work will support nested conditionals by adding multiple flag species for multiple comparisons.

An important direction for future research concerns reducing the error in our construction, and understanding how it builds up over time. We noticed that different algorithms, even computing the same function, accumulate varying levels of error. For example, as shown in the full version of this paper, the error of the Sub module increases with the magnitude of the operands, and also increases the closer they are. However, we also found an alternative way to subtract, that keeps the error constant and independent of the operands (see the full version of this paper) at the cost of a slower run-time.

Our error analysis shows that for most examples we tried, but not all, error builds up over the course of the computation. For *CRN++* programs where the error builds up in this way, there is some maximum input complexity beyond which the error overwhelms the output. Can all *CRN++* programs be refactored (preferably automatically) to bound the cumulative error of every module such that it does not build up over time? Note that if this were possible, we would obtain another, more efficient, way to achieve Turing universality.

To the best of our knowledge we are the first to provide an imperative programming language which compiles to chemical reaction networks. Moreover, we build tools that can help users get a better understanding of CRNs and improve their design. Although absolutely correct computation is not achieved, we provide tools that help understand why error occurs, and thus help improve the design of CRNs. We release our toolkit as open-source, to encourage new research and improvement of the *CRN++*, with the hope of advancing the engineering of information processing molecular systems.

Acknowledgment. We thank the fellow students of EE 381V (Programming With Molecules) at The University of Texas at Austin for constructive discussions on the material presented in this paper. We also thank Keenan Breik, Cameron Chalk, Milos Gligoric, Aleksandar Milicevic, Boya Wang and Kaiyuan Wang for their feedback on this work. This research was partially supported by the US National Science Foundation under Grants Nos. CCF-1618895, CCF-1718903, CCF-1652824, and CCF-1704790.

A Modules

Table 1. CRN++ Modules. The first column denotes the type of the module. The restrictions column imposes compile-time restrictions for using modules, here \neq is used to mean different species (not values). The output column shows the value of outputs at the steady state. Finally, the CRN column shows chemical reactions implementing the module.

Type	Restrictions	Output (Steady State)	CRN
$ld[A,B]$	$B \neq A$	$B := A$	$A \longrightarrow A + B$ $B \longrightarrow \emptyset$
$add[A,B,C]$	$C \neq A \wedge C \neq B$	$C := A + B$	$A \longrightarrow A + C$ $B \longrightarrow B + C$ $C \longrightarrow \emptyset$
$sub[A,B,C]$	$C \neq A \wedge C \neq B$	$C := \begin{cases} A - B, & A > B \\ 0, & \text{otherwise} \end{cases}$	$A \longrightarrow A + C$ $B \longrightarrow B + H$ $C \longrightarrow \emptyset$ $C + H \longrightarrow \emptyset$
$mul[A,B,C]$	$C \neq A \wedge C \neq B$	$C := A \cdot B$	$A + B \longrightarrow A + B + C$ $C \longrightarrow \emptyset$
$div[A,B,C]$	$C \neq A \wedge C \neq B$	$C := A/B$	$A \longrightarrow A + C$ $B + C \longrightarrow B$
$sqrt[A,B]$	$B \neq A$	$B := \sqrt{A}$	$A \xrightarrow{1} A + B$ $B + B \xrightarrow{\frac{1}{2}} \emptyset$
$am[A,B]$	$A \neq B$	$A := \begin{cases} A + B, & A > B \\ 0, & B > A \end{cases}$ $B := \begin{cases} 0, & A > B \\ A + B, & B > A \end{cases}$	$A + B \longrightarrow A + T$ $B + A \longrightarrow B + T$ $T + A \longrightarrow A + A$ $T + B \longrightarrow B + B$
$cmp[A,B]$	$A \neq B$	Sets flag species	* Two CRNs (mapping and AM) triggering in two consecutive phases (as discussed in the Technique)

References

1. *CRN++* Github page. <https://github.com/marko-vasic/crnPlusPlus>
2. Angluin, D., Aspnes, J., Eisenstat, D.: A simple population protocol for fast robust approximate majority. *Distrib. Comput.* **21**(2), 87–102 (2008)
3. Baccouche, A., Montagne, K., Padirac, A., Fujii, T., Rondelez, Y.: Dynamic DNA-toolbox reaction circuits: a walkthrough. *Methods* **67**(2), 234–249 (2014)
4. Bournez, O., Graça, D.S., Pouly, A.: Polynomial time corresponds to solutions of polynomial ordinary differential equations of polynomial length. *J. ACM* **64**(6), 38 (2017)
5. Buisman, H.J., ten Eikelder, H.M.M., Hilbers, P.A.J., Liekens, A.M.L.: Computing algebraic functions with biochemical reaction networks. *Artif. Life* 5–19 (2009)
6. Cardelli, L., Csikász-Nagy, A.: The cell cycle switch computes approximate majority. *Sci. Rep.* **2**, 656 (2012)
7. Chen, Y.J., et al.: Programmable chemical controllers made from DNA. *Nat. Nanotechnol.* **8**(10), 755 (2013)
8. CRNSimulator Mathematica package. <http://users.ece.utexas.edu/~soloveichik/crn simulator.html>
9. Fages, F., Le Guludec, G., Bournez, O., Pouly, A.: Strong turing completeness of continuous chemical reaction networks and compilation of mixed analog-digital programs. In: Feret, J., Koepl, H. (eds.) *CMSB 2017. LNCS*, vol. 10545, pp. 108–127. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-67471-1_7
10. Huang, D.A., Jiang, J.H.R., Huang, R.Y., Cheng, C.Y.: Compiling program control flows into biochemical reactions. In: *Proceedings of the International Conference on Computer-Aided Design*, pp. 361–368 (2012)
11. Jiang, H., Riedel, M., Parhi, K.: Synchronous sequential computation with molecular reactions. In: *2011 48th ACM/EDAC/IEEE Design Automation Conference (DAC)*, pp. 836–841 (2011)
12. Lachmann, M., Sella, G.: The computationally complete ant colony: global coordination in a system with no hierarchy. In: Morán, F., Moreno, A., Merelo, J.J., Chacón, P. (eds.) *ECAL 1995. LNCS*, vol. 929, pp. 784–800. Springer, Heidelberg (1995). https://doi.org/10.1007/3-540-59496-5_343
13. Magasco, M.O.: Chemical kinetics is Turing universal. *Phys. Rev. Lett.* **78**(6), 1190 (1997)
14. Salehi, S.A., Parhi, K.K., Riedel, M.D.: Chemical reaction networks for computing polynomials. *ACS Synth. Biol.* **6**(1), 76–83 (2017)
15. Soloveichik, D., Seelig, G., Winfree, E.: DNA as a universal substrate for chemical kinetics. *Proc. Natl. Acad. Sci.* **107**(12), 5393–5398 (2010)