



Programming
Cryptol

Cryptol:

The Language of Cryptography



Programming Cryptol

*Levent Erkök
Dylan McNamee
Joe Kiniry
Iavor Diatchki
John Launchbury*

Galois, Inc.
421 SW Sixth Avenue, Suite 300
Portland, OR 97204

Acknowledgements

Cryptol has been under development for over a decade with many people contributing to its design and implementation. Those people include (but are not limited to) Aaron Tomb, Adam Wick, Brian Huffman, Matt Sottile, Adam Foltzer, Joe Hendrix, Trevor Elliott, Lee Pike, Mark Tullsen, David Lazar, Joel Stanley, Jeff Lewis, Andy Gill, Edward Yang, Ledah Casburn, Jim Teisher, Sigbjørn Finne, Mark Shields, Philip Weaver, Magnus Carlsson, Fergus Henderson, Joe Hurd, Thomas Nordin, John Matthews and Sally Browning. In addition, much of the work on Cryptol has been funded by, and lots of design input was provided by the team at the NSA's Trusted Systems Research Group, including Brad Martin, Frank Taylor and Sean Weaver.

IMPORTANT NOTICE

This documentation is furnished for informational use only and is subject to change without notice. Galois, Inc. assumes no responsibility or liability for any errors or inaccuracies that may appear in this documentation.

Copyright 2003–2014 Galois, Inc. All rights reserved by Galois, Inc.

This documentation, or any part of it, may not be reproduced, displayed, photocopied, transmitted, or otherwise copied in any form or by any means now known or later developed, such as electronic, optical, or mechanical means, without express written authorization from Galois, Inc.

Warning: This computer program and its documentation are protected by copyright law and international treaties. Any unauthorized copying or distribution of this program, its documentation, or any portion of it, may result in severe civil and criminal penalties, and will be prosecuted under the maximum extent possible under the law.

The software installed in accordance with this documentation is copyrighted and licensed by Galois, Inc. under separate license agreement. This software may only be used pursuant to the terms and conditions of such license agreement. Any other use may be a violation of law.

TRADEMARKS

Cryptol is a registered trademark of Galois, Inc. in the United States and other countries. UNIX is a registered trademark of The Open Group in the U. S. and other countries. Linux is a registered trademark of Linus Torvalds.

Other company or product names mentioned herein may be trademarks or registered trademarks of their respective owners. Trademark specifications are subject to change without notice. All terms mentioned in this documentation that are known to be trademarks or service marks have been appropriately capitalized to the best of our knowledge; however, Galois cannot attest to the accuracy of all trademark information. Use of a term in this documentation should not be regarded as affecting the validity of any trademark or service mark.

Galois, Inc.
421 SW Sixth Avenue, Suite 300
Portland, OR 97204

Contents

Contents	v
Preface	1
1 Installation and Tool Use	3
1.1 Getting started	3
1.2 Technicalities	4
1.2.1 Language features	4
1.2.2 Commands	5
2 A Crash Course in Cryptol	9
2.1 Basic data types	9
2.2 Bits: Booleans	9
2.3 Words: Numbers	10
2.4 Tuples: Heterogeneous collections	11
2.5 Sequences: Homogeneous collections	11
2.5.1 Enumerations	12
2.5.2 Comprehensions	12
2.5.3 Appending and indexing	13
2.5.4 Finite and infinite sequences	14
2.5.5 Manipulating sequences	14
2.5.6 Shifts and rotates	16
2.6 Words revisited	16
2.7 Characters and strings	17
2.8 Records: Named collections	18
2.9 The zero	19
2.10 Arithmetic	20
2.11 Types	22
2.11.1 Monomorphic types	22
2.11.2 Polymorphic types	23
2.11.3 Predicates	25
2.11.4 Why typed?	25
2.12 Defining functions	26
2.12.1 Local names: <code>where</code> -clauses	27
2.12.2 λ -expressions	27

2.12.3	Using zero in functions	28
2.13	Recursion and recurrences	28
2.14	Stream equations	31
2.15	Type synonyms	32
2.16	Type classes	33
2.17	Type vs. value variables	33
2.17.1	Positional vs. named type arguments	34
2.17.2	Type context vs. variable context	35
2.17.3	Inline argument type declarations	35
2.18	The road ahead	35
3	Classic ciphers	37
3.1	Caesar’s cipher	38
3.2	Vigenère cipher	38
3.3	The atbash	39
3.4	Substitution ciphers	40
3.5	The scytale	41
4	The Enigma machine	43
4.1	The plugboard	43
4.2	Scrambler rotors	44
4.3	Connecting the rotors: notches in action	45
4.4	The reflector	47
4.5	Putting the pieces together	47
4.6	The state of the machine	48
4.7	Encryption and decryption	49
5	High-assurance programming	51
5.1	Writing properties	51
5.1.1	Property-function correspondence	52
5.1.2	Capturing test vectors	53
5.1.3	Polymorphic properties	53
5.2	Establishing correctness	54
5.2.1	Formal proofs	54
5.2.2	Counterexamples	55
5.2.3	Dealing with polymorphism	55
5.2.4	Conditional proofs	56
5.3	Automated random testing	57
5.4	Checking satisfiability	58
6	AES: The Advanced Encryption Standard	61
6.1	Parameters	62
6.2	Polynomials in $\text{GF}(2^8)$	62
6.3	The <code>subBytes</code> transformation	64
6.4	The <code>shiftRows</code> transformation	66
6.5	The <code>mixColumns</code> transformation	67
6.6	Key expansion	68

Contents

6.7	The <code>AddRoundKey</code> transformation	71
6.8	AES encryption	71
6.9	Decryption	73
6.9.1	The <code>InvSubBytes</code> transformation	73
6.9.2	The <code>InvShiftRows</code> transformation	74
6.9.3	The <code>InvMixColumns</code> transformation	74
6.10	The inverse cipher	74
6.11	Correctness	75
A	Solutions to selected exercises	77
1.1	Getting started	77
2.2	Bits: Booleans	77
2.3	Words: Numbers	78
2.4	Tuples: Heterogeneous collections	78
2.5	Sequences: Homogeneous collections	78
2.6	Words revisited	83
2.8	Records: Named collections	84
2.9	The zero	85
2.10	Arithmetic	85
2.11	Types	87
2.12	Defining functions	88
2.13	Recursion and recurrences	90
2.14	Stream equations	92
2.15	Type synonyms	92
3.1	Caesar’s cipher	93
3.2	Vigenère cipher	95
3.3	The atbash	96
3.4	Substitution ciphers	96
3.5	The scytale	97
4.1	The plugboard	98
4.2	Scrambler rotors	98
4.3	Connecting the rotors: notches in action	99
4.4	The reflector	100
4.5	Putting the pieces together	101
4.7	Encryption and decryption	102
5.1	Writing properties	102
5.2	Establishing correctness	103
5.3	Automated random testing	105
5.4	Checking satisfiability	106
6.2	Polynomials in $GF(2^8)$	107
6.3	The <code>SubBytes</code> transformation	108
6.4	The <code>ShiftRows</code> transformation	110
6.5	The <code>MixColumns</code> transformation	110
6.6	Key expansion	111
6.8	AES encryption	112
6.9	Decryption	112

B Cryptol primitive functions	115
C Enigma simulator	117
D AES in Cryptol	121
E Cryptol Syntax	127
E.1 Layout	127
E.2 Comments	127
E.3 Identifiers	128
E.4 Keywords and Built-in Operators	128
E.5 Numeric Literals	129
E.6 Bits	129
E.7 If Then Else with Multiway	130
E.8 Tuples and Records	130
E.9 Sequences	131
E.10 Functions	132
E.11 Local Declarations	132
E.12 Explicit Type Instantiation	132
E.13 Demoting Numeric Types to Values	132
E.14 Explicit Type Annotations	133
E.15 Type Signatures	133
E.16 Type Synonym Declarations	133
F The Cryptol Grammar	135
Glossary	137
Bibliography	139
Index	141

Preface

Cryptol is a domain-specific language for programming, executing, testing, and formally reasoning about streams of bits. It particularly excels at specifying and reasoning about cryptographic algorithms. It has been designed to reduce the gap that currently exists between the *specification* of a cryptographic algorithm and its executable *implementation*. As a result, a well-written Cryptol program will look very much like the specification of the algorithm it implements, and is also executable.

Our goal in this book is to both teach you the Cryptol language and provide a reference text for the use of the Cryptol system. The Cryptol system provides: (1) a REPL for experimentation, (2) a parser and typechecker for Cryptol programs, (3) an interpreter for executing Cryptol programs, (4) automatically checking the correctness of Cryptol programs via randomized testing, and (5) formally verifying Cryptol programs through the use of an SMT solver.

We demonstrate Cryptol in action via normal programming problems, traditional cryptographic techniques (such as substitution ciphers), historical cryptographic mechanisms (such as the Enigma), and modern algorithms (such as DES, SHA, and AES), focusing on how they are elegantly modeled using the Cryptol language. Since our emphasis is on programming, we introduce some of the techniques that are useful for the working programmer, including the use of Cryptol's validation and verification tools that directly support high-assurance programming.

Our approach is hands on. We strongly urge the reader to try out the material using the Cryptol toolset directly. It would be most effective to read this document while trying the code and interacting with Cryptol itself. For information on how to download and use the Cryptol toolset, please visit www.cryptol.net.

Most exercises come with solutions provided at the end. We urge the readers to refer to solutions as a last resort as they work through the exercises.

Note. This book covers Cryptol version 2, which is still under development. It is possible that the syntax and semantics for the language or commands will change. You may encounter bugs or unexpected behaviors. Please let us know if you encounter any problems, by sending email to cryptol@galois.com or by visiting our GitHub presence linked from the Cryptol website.

The authors of this book are Levent Erkök, Dylan McNamee, Joseph Kiniry, Iavor Diatchki, and John Launchbury, with contributions from Magnus Carlsson, Kyle Carter, Trevor Elliott, Adam Foltzer, Brian Huffman, David Lazar, John Matthews, Eric Mertens, Matt Sottile, Aaron Tomb, and Adam Wick.

Chapter 1

Installation and Tool Use

The best way to learn Cryptol is to: (a) read this book and related course notes and, more importantly, (b) attempt the dozens of problems included herein. The only way to really learn a new language is to use it, so installing and using the Cryptol tool is critical to your success in using Cryptol.

1.1 Getting started

How you download Cryptol and install it on your system is platform specific [4]. Once installed, we use Cryptol from inside a terminal window, by interacting with its read-eval-print loop. On a Linux/Mac machine, this simply amounts to typing:

```
$ cryptol
```



```
Loading module Cryptol  
Cryptol>
```

Of course, your version number may be different, but for this document we will assume you are running at least version 2.0¹.

On Windows, you typically click on the desktop icon to run Cryptol in a command window. Otherwise, the interaction mode is exactly the same, regardless of which platform you use to run Cryptol.

Once you have the `cryptol>` prompt displayed, you are good to go. You can directly type expressions (not declarations) and have Cryptol evaluate them for you. The extension for Cryptol program files is `.cry`. If you place your program in a file called `prog.cry`, then you can load it into Cryptol like this:

```
Cryptol> :l prog.cry
```

¹Version 2.0 of cryptol is a significant change from version 1. If you are already familiar with Cryptol version 1, there is a document in the Cryptol release that summarizes the changes.

or, by calling Cryptol from your shell as follows:

```
$ cryptol prog.cry
```

Exercise 1. Obtain and install Cryptol on your machine. Start it up and type:

```
Cryptol> "Hello World!"
```

What do you see printed?

Exercise 2. Try the above exercise, after first issuing the following command:

```
Cryptol> :set ascii=on
```

Why do you think this is not the default behavior?

1.2 Technicalities

Before we dive into various aspects of Cryptol, it is good to get some of the technicalities out of the way. Feel free to skim through these items for future reference. The summary below describes language features, as well as commands that are available at the `cryptol>` prompt. Commands all begin with the `:` character.

1.2.1 Language features

The Cryptol language is a size-polymorphic dependently-typed programming language with support for polymorphic recursive functions. It has a small syntax tuned for applied cryptography, a lightweight module system, a pseudo-Real/Eval/Print/Loop (REPL) top-level, and a rich set of built-in tools for performing high-assurance (iterate) programming. Cryptol performs fairly advanced type inference, though as with most mainstream strongly typed functional languages, types can be manually specified as well. What follows is a brief tour of Cryptol's most salient language features.

Case sensitivity Cryptol identifiers are case sensitive. `A` and `a` are two different things.

Indentation and whitespace Cryptol uses indentation-level (instead of `{}`'s) to denote blocks. Whitespace within a line is immaterial, as is the specific amount of indentation. However, consistent indentation will save you tons of trouble down the road! Do not mix tabs and spaces for your indentation. Spaces are generally preferred.

Escape characters Long lines can be continued with the end-of-line escape character `\`, as in many programming languages. There are no built-in character escape characters, as Cryptol performs no interpretation on bytes beyond printing byte streams out in ASCII, as discussed above.

Comments Block comments are enclosed in `/*` and `*/`, and they can be nested. Line comments start with `//` and run to the end of the line.

Order of definitions The order of definitions is immaterial. You can write your definitions in any order, and earlier entries can refer to latter ones.

1.2. Technicalities

Typing Cryptol is strongly typed. This means that the interpreter will catch most common mistakes in programming during the type-checking phase, before runtime.

Type inference Cryptol has type inference. This means that the user can omit type signatures because the inference engine will supply them.

Type signatures While writing type signatures are optional, writing them down is considered good practice.

Polymorphism Cryptol functions can be polymorphic, which means they can operate on many different types. Beware that the type which Cryptol infers might be too polymorphic, so it is good practice to write your signatures, or at least check what Cryptol inferred is what you had in mind.

Module system Each Cryptol file defines a *module*. Modules allow Cryptol developers to manage which definitions are exported (the default behavior) and which definitions are internal-only (*private*). At the beginning of each Cryptol file, you specify its name and use `import` to specify the modules on which it relies. Definitions are `public` by default, but you can hide them from modules that import your code via the `private` keyword at the start of each private definition, like this:

```
module test where
private
  hiddenConst = 0x5      // hidden from importing modules
// end of indented block indicates symbols are available to importing modules
revealedConst = 0x15
```

Note that the filename should correspond to the module name, so `module test` must be defined in a file called `test.cry`.

Literate programming You can feed \LaTeX files to Cryptol (i.e., files with extension `.tex`). Cryptol will look for `\begin{code}` and `\end{code}` marks to extract Cryptol code. Everything else will be comments as far as Cryptol is concerned. In fact, the book you are reading is a Literate Cryptol program.

Completion On UNIX-based machines, you can press `tab` at any time and Cryptol will suggest completions based on the context. You can retrieve your prior commands using the usual means (arrow keys or Emacs keybindings).

1.2.2 Commands

Querying types You can ask Cryptol to tell you the type of an expression by typing `:type <expr>` (or `:t` for short). If `f00` is the name of a definition (function or otherwise), you can ask its type by issuing `:type f00`. It is common practice to define a function, ask Cryptol its type, and copy the response back to your source code. While this is somewhat contrived, it is usually better than not writing signatures at all. In order to query the type of an infix operator (e.g., `+`, `==`, etc.) you will need to surround the operator with `()`'s, like this:

```
Cryptol> :t (+)
+ : {a} (Arith a) => a -> a -> a
```

Browsing definitions The command `:browse` (or `:b` for short) will display all the names you have defined, along with their types.

Getting help The command `:help` will show you all the available commands. Other useful implicit help invocations are: (a) to type `tab` at the `Cryptol>` prompt, which will list all of the operators available in Cryptol code, (b) typing `:set` with no argument, which shows you the parameters that can be set, and (c), as noted elsewhere, `:browse` to see the names of functions and type aliases you have defined, along with their types.

Option	Default value	Meaning
<code>ascii</code>	<code>off</code>	print sequences of bytes as a string
<code>base</code>	<code>10</code>	numeric base for printing words
<code>debug</code>	<code>off</code>	whether to print verbose debugging information
<code>infLength</code>	<code>5</code>	number of elements to show from an infinite sequence
<code>iteSolver</code>	<code>off</code>	whether <code>:prove</code> calls out for help on recursive functions
<code>prover</code>	<code>cvc4</code>	which SMT solver to use for <code>:prove</code>
<code>tests</code>	<code>100</code>	number of tests to run for <code>:check</code>
<code>warnDefaulting</code>	<code>on</code>	

Environment options A variety of environment options are set through the use of the `:set` command. These options may change over time and some options may be available only on specific platforms. The current options are summarized in section 1.2.2.

Quitting You can quit Cryptol by using the command `:quit` (aka `:q`). On Mac/Linux you can press `Ctrl-D`, and on Windows use `Ctrl-Z`, for the same effect.

Loading and reloading files You load your program in Cryptol using `:load <filename>` (or `:l` for short). However, it is customary to use the extension `.cry` for Cryptol programs. If you edit the source file loaded into Cryptol from a separate context, you can reload it into Cryptol using the command `:reload` (abbreviated `:r`).

Invoking your editor You can invoke your editor using the command `:edit` (abbreviated `:e`). The default editor invoked is `vi`. You override the default using the standard `EDITOR` environmental variable in your shell.

Running shell commands You can run Unix shell commands from within Cryptol like this: `:! cat test.cry`.

Changing working directory You can change the current working directory of Cryptol like this: `:cd some/path`. Note that the path syntax is platform-dependent.

Loading a module At the Cryptol prompt you can load a module by name with the `:module` command.

The next three commands all operate on *properties*. All take either one or zero arguments. If one argument is provided, then that property is the focus of the command; otherwise all properties in the current context are checked. All three commands are covered in detail in chapter 5.

1.2. Technicalities

Checking a property through random testing The `:check` command performs random value testing on a property to increase one’s confidence that the property is valid. See section 5.3 for more detailed information.

Verifying a property through automated theorem proving The `:prove` command uses an external SMT solver to attempt to automatically formally prove that a given property is valid. See subsection 5.2.1 for more detailed information.

Finding a satisfying assignment for a property The `:sat` command uses an external SAT solver to attempt to find a satisfying assignment to a property. See section 5.4 for more detailed information.

Type specialization Discuss `:debug_specialize`.

The next chapter provides a “crash course” introduction to the Cryptol programming language.

Chapter 2

A Crash Course in Cryptol

Before we can delve into cryptography, we have to get familiar with Cryptol. This chapter provides an introduction to Cryptol, just to get you started. The exposition is not meant to be comprehensive, but rather as an overview to give you a feel of the most important tools available. If a particular topic appears hard to approach, feel free to skim it over for future reference.

A full language reference is beyond the scope of this document at this time.

The full grammar for Cryptol is included in Appendix F.

2.1 Basic data types

Cryptol provides four basic data types: bits, sequences, tuples, and records. Words (i.e., numbers) are a special case of sequences. Note that, aside from bits, all other Cryptol types can be nested as deep as you like. That is, we can have records of sequences containing tuples that comprise of other records, etc., giving us a rich type-system for precisely describing the shapes of data our programs manipulate.

While Cryptol is statically typed, it uses type inference to supply unspecified types. That is, the user does *not* have to write the types of all expressions, they will be automatically inferred by the type-inference engine. Of course, in certain contexts the user might choose to supply a type explicitly. The notation is simple: we simply put the expression, followed by `:` and the type. For instance:

```
12 : [8]
```

means the value 12 has type `[8]`, i.e., it is an 8-bit word. We shall see other examples of this in the following discussion.

2.2 Bits: Booleans

The type `Bit` represents a single bit of information. There are precisely two values of this type: `True` and `False`. Bit values play an important role in Cryptol, as we shall see in detail shortly. In particular, the test expression in an `if-then-else` statement must have the type `Bit`. The logical operators `&&` (and), `||` (or), `^` (xor), and `~` (complement) provide the basic operators that act on bit values.

Exercise 1. Type in the following expressions at the Cryptol prompt, and observe the output:

```

True
false
False : Bit
if True && False then 3 else 4
False || True
(True && False) ^ True
~False
~(False || True)

```

Remember that Cryptol is case sensitive, and hence `false` is different from `False`.

Tip. Cryptol provides extensive command line/tab completion; use up/down-arrow to see your previous commands, hit tab to complete identifier names, etc.

2.3 Words: Numbers

A word is simply a numeric value, corresponding to the usual notion of numbers. To match our observation of how cryptographers use numbers, Cryptol only supports non-negative (≥ 0) integer values (i.e., no floating point or negative numbers). However, numbers can be arbitrarily large: There is no predefined maximum that we are limited to. By default, Cryptol prints numbers in base 16. You might find it useful to set the output base to be 10 while working on the following example. To do so, use the command:

```
:set base=10
```

The most common values for this setting are 2 (binary), 8 (octal), 10 (decimal), and 16 (hexadecimal). Conversely, we can *write* numbers in these bases in Cryptol programs too:

```

0b11111011110    // binary
0o3736           // octal
2014             // decimal
0x7de           // hexadecimal

```

For printing values in arbitrary bases, Cryptol uses the notation `0<base>digits`, where `base` is the base value and the `digits` are the numeric value in that particular base. E.g., the above value is equal to `0<7>5605` and `0<20>50e`. One cannot input a value in a non-standard base.¹

Note. Decimal numbers pose a problem in a bit-precise language like Cryptol. Numbers represented in a base that is a power of two unambiguously specify the number of bits required to store each digit. For example `0b101` takes three bits to store. A hexadecimal digit takes 4 bits to store, so `0xabc` needs 12 bits. On the other hand, in decimal, the number of bits is ambiguous. A decimal digit could require anywhere from 1 to 4 bits to represent. When given a choice, Cryptol assumes the *smallest* number of bits required to represent a decimal number. This is why Cryptol often prints messages like `Assuming a = 3`; the value emitted are the number of bits necessary to faithfully represent the decimal value on the corresponding line.

Exercise 2. Experiment with different output bases by issuing `:set base=10`, and other base values. Also try writing numbers directly in different bases at the command line, such as `0o1237`. Feel free to try other

¹Cryptol does not support the input of numbers in arbitrary bases—the use of non-standard bases (i.e., beyond base 2, 8, 10, and 16) is vanishingly rare and thus not worth the trouble in complicating the Cryptol parser.

2.4. Tuples: Heterogeneous collections

bases. What is the hexadecimal version of the octal number `0o7756677263570015`? Observe that `:base=` can be set to anything between 2 and 36. Why does Cryptol stop at 36?

Note. We will revisit the notion of numbers in Section 2.6, after we learn about sequences.

2.4 Tuples: Heterogeneous collections

A tuple is a simple collection of arbitrary ordered values of arbitrary types, written in parentheses. A tuple is at least a pair; it has at least two elements², and can be arbitrarily nested with other types. Elements are comma separated.

Two tuples are equal in the standard fashion: if they have the same arity, their types are pairwise comparable (see section 2.16), and their values are pairwise identical. Note that their types need not be pairwise identical; for example, consider the expression `('A', 0) == (65, 0)`. The type of the LHS is `a (fin a) => ([8], [a])` while the RHS is `a, b (a >= 7, fin a, fin b) => ([a], [b])`.

Exercise 3. Try out the following tuples:

```
(1, 2+4)
(True, False, True ^ False)
((1, 2), False, (3-1, (4, True)))
```

Projecting values from tuples Use a `.` followed by n to project the $n + 1$ -th component of a tuple. Nested projection is not supported at this time.

Exercise 4. Try out the following examples:

```
(1, 2+4).0
(1, 2+4).1
((1, 2), False, (3-1, (4, True))).2
```

Write a projection to extract the value `False` from the expression:

```
((1, 2), (2, (4, True), 6), False)
```

Tip. While projections can come in handy, we rarely see them used in Cryptol programs. As we shall see later, Cryptol's powerful pattern-matching mechanism provides a much nicer and usable alternative for extracting parts of tuples and other composite data values.

2.5 Sequences: Homogeneous collections

While tuples contain heterogeneous data, sequences are used for homogeneous collections of values, akin to value arrays in more traditional languages. A sequence contains elements of any *single* type, even sequences themselves, arbitrarily nested. We simply write a sequence by enclosing it within square brackets with comma-separated elements.

Exercise 5. Try out the following sequences:

²Tuples with zero and one element are part of the underlying mathematics of Cryptol's tuple theory but are not supported in its concrete syntax because doing so unnecessarily complicates the parser and program comprehension.

```
[1, 2]
[[1, 2, 3], [4, 5, 6], [7, 8, 9]]
```

Note how the latter example can be used as the representation of a 3×3 matrix.

Tip. The most important thing to remember about a sequence is that its elements must be of exactly the same type.

Exercise 6. Type in the following expressions to Cryptol and observe the type-errors:

```
[True, [True]]
[[1, 2, 3], [4, 5]]
```

2.5.1 Enumerations

Cryptol enumerations allow us to write sequences more compactly, instead of listing the elements individually. An enumeration is a means of writing a sequence by providing a (possibly infinite) range. Cryptol enumerations are not equivalent to mainstream programming languages' notions of enumeration types, other than both kinds of constructs guarantee that enumeration elements are distinct.

Exercise 7. Explore various ways of constructing enumerations in Cryptol, by using the following expressions:

```
[1 .. 10]           // increment with step 1
[1, 3 .. 10]        // increment with step 2 (= 3-1)
[10, 9 .. 1]        // decrement with step 1 (= 10-9)
[10, 9 .. 20]       // decrement with step 1 (= 10-9)
[10, 7 .. 1]        // decrement with step 3 (= 10-7)
[10, 11 .. 1]       // increment with step 1
```

2.5.2 Comprehensions

A Cryptol comprehension is a way of programmatically computing the elements of a new sequence, out of the elements of existing ones. The syntax is reminiscent of the set comprehension notation from ordinary mathematics, generalized to cover parallel branches (as explained in the exercises below). Note that Cryptol comprehensions are not generalized numeric comprehensions (like summation, product, maximum, or minimum), though such comprehensions can certainly be defined using Cryptol comprehensions.

Exercise 8. A comprehension with a single arm is called a *cartesian comprehension*. We can have one or more components in a cartesian comprehension. Experiment with the following expressions:

```
[ (x, y) | x <- [1 .. 3], y <- [4, 5] ]
[ x + y | x <- [1 .. 3], y <- [] ]
[ (x + y, z) | x <- [1, 2], y <- [1], z <- [3, 4] ]
```

What is the number of elements in the resulting sequence, with respect to the sizes of components?

Note. Recall that, when you type the expressions above, you will get messages from Cryptol such as `Assuming a = 2`. This is Cryptol letting you know it has decided to use 2 bits to represent, for example, the value 3 in `[1 .. 3]`. This information may not seem to matter now but it can be very helpful later on.

2.5. Sequences: Homogeneous collections

Exercise 9. A comprehension with multiple arms is called a *parallel comprehension*. We can have any number of parallel arms. The contents of each arm will be *zipped* to obtain the results. Experiment with the following expressions:

```
[ (x, y) | x <- [1 .. 3] | y <- [4, 5] ]
[ x + y | x <- [1 .. 3] | y <- [] ]
[ (x + y, z) | x <- [1, 2] | y <- [1] | z <- [3, 4] ]
```

What is the number of elements in the resulting sequence, with respect to the sizes of the parallel branches?

Tip. One can mix parallel and cartesian comprehensions, where each parallel arm can contain multiple cartesian generators, or vice-versa.

Tip. While Cryptol comprehensions *look* like standard mathematical comprehensions, one must remember that the codomain of Cryptol comprehensions is a sequence type of some kind, *not* a set.

Comprehensions may be nested. In this pattern, the element value expression of the outer nesting is a sequence comprehension (which may refer to values generated by the outer generator). The pattern looks like

```
      [ [ <expr with x & y> // \
        | y <- [1 .. 5] // inner generator -- outer elements
this: ] /
      | x <- [1 .. 5] // outer generator
      ]
```

Exercise 10. Use a nested comprehension to write an expression to produce a 3×3 matrix (as a sequence of sequences), such that the ij th entry contains the value (i, j) .

2.5.3 Appending and indexing

For sequences, the two basic operations are appending (`#`) and selecting elements out (`@`, `@@`, `!`, and `!!`). Forward selection operator (`@`), starts counting from the beginning, while the backward selection operator (`!`) starts from the end. Indexing always starts at zero: That is `xs @ 0` is the first element of `xs`, while `xs ! 0` is the last. The permutation versions (`@@` and `!!`, respectively) allow us to concisely select multiple elements: they allow us to extract elements in any order (which makes them very useful for permuting sequences).

Exercise 11. Try out the following Cryptol expressions:

```
[ ] # [1, 2]
[1, 2] # [ ]
[1 .. 5] # [3, 6, 8]
[0 .. 9] @ 0
[0 .. 9] @ 5
[0 .. 9] @ 10
[0 .. 9] @@ [3, 4]
[0 .. 9] @@ [ ]
[0 .. 9] @@ [9, 12]
[0 .. 9] @@ [9, 8 .. 0]
[0 .. 9] ! 0
[0 .. 9] ! 3
[0 .. 9] !! [3, 6]
```

```
[0 .. 9] !! [0 .. 9]
[0 .. 9] ! 12
```

Exercise 12. The permutation operators (`@@` and `!!`) can be defined using sequence comprehensions. Write an expression that selects the even indexed elements out of the sequence `[0 .. 10]` first using `@@`, and then using a sequence comprehension.

2.5.4 Finite and infinite sequences

So far we have only seen finite sequences. An infinite sequence is one that has an infinite number of elements, corresponding to streams. Cryptol supports infinite sequences, where the elements are accessed *on-demand*. This implies that Cryptol will *not* go into an infinite loop just because you have created an infinite sequence: it will lazily construct the sequence and make its elements available as demanded by the program.

Exercise 13. Try the following infinite enumerations:

```
[1:[32] ... ]
[1:[32], 3 ...]
[1:[32] ...] @ 2000
[1:[32], 3 ...] @@ [300, 500, 700]
[100, 102 ...]
```

Note. Note that we are explicitly telling Cryptol to use 32-bit words as the elements. The reason for doing so will become clear when we study arithmetic shortly.

Exercise 14. What happens if you use the reverse index operator (`!`) on an infinite sequence? Why?

2.5.5 Manipulating sequences

Sequences are at the heart of Cryptol, and there are a number of built-in functions for manipulating them in various ways. It is worthwhile to try the following exercises to gain basic familiarity with the basic operations.

Exercise 15. Try the following expressions:

```
take' {3} [1 .. 12]
drop' {3} [1 .. 12]
split' {3} [1 .. 12]
groupBy' {3} [1 .. 12]
join [[1 .. 4], [5 .. 8], [9 .. 12]]
join [[1, 2, 3], [4, 5, 6], [7, 8, 9], [10, 11, 12]]
transpose [[1, 2, 3, 4], [5, 6, 7, 8]]
transpose [[1, 2, 3], [4, 5, 6], [7, 8, 9]]
```

And for fun, think about what these should produce:

```
join [1,1]
transpose [1,2]
```

2.5. Sequences: Homogeneous collections

Exercise 16. Based on your intuitions from the previous exercise, derive laws between the following pairs of functions: `take` and `drop`; `join` and `split`; `join` and `groupBy`; `split` and `groupBy` and `transpose` and itself. For instance, `take` and `drop` satisfy the following equality:

```
(take' {n} xs) # (drop' {n} xs) == xs
```

whenever `n` is between 0 and the length of the sequence `xs`. Note that there might be multiple laws these functions satisfy.

Exercise 17. What is the relationship between the append operator `#` and `join`?

Type-directed splits We have studied the functions `groupBy` and `splitBy` above. Cryptol also provides a function `split` that can split a sequence into any number of equal-length segments. A common way to use `split` is to be explicit about the type of its result, instead of passing arguments as we did above with `splitBy` and `groupBy`.

```
Cryptol> split [1..12] : [1][12][8]
[[1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12]]
Cryptol> split [1..12] : [2][6][8]
[[1, 2, 3, 4, 5, 6], [7, 8, 9, 10, 11, 12]]
Cryptol> split [1..12] : [3][4][8]
[[1, 2, 3, 4], [5, 6, 7, 8], [9, 10, 11, 12]]
```

Here is what happens if we do *not* give an explicit signature on the result:

```
Cryptol> split [1..12]
<polymorphic value>
Cryptol> :t split [1..12]
split [1 .. 12] : {a, b, c} (a >= 4, fin a, fin c,
                        12 == b * c) => [b][c][a]
```

A complex type signature like this one first defines a set of type variables `{a, b, c}`, a set of constraints on those variables (`a >= 4, fin a, fin c, 12 == b * c`), `a =>` and finally the shape description. In this case, Cryptol's `[b][c][a]` is telling us that the result will be a sequence of `b` things, each of which is a sequence of `c` things, each of which is a word of size `a`. The type constraints tell us that `a` is at least 4, because the maximum element of the sequence is 12, and it takes at least 4 bits to represent the value 12. The constraints are that `b * c == 12`, which means we should completely cover the entire input, and that the lengths `a` and `c` need to be finite. As you can see, `split` is a very powerful function. The flexibility afforded by `split` comes in very handy in Cryptol. We shall see one example of its usage later in Section 3.5.

Exercise 18. With a sequence of length 12, as in the above example, there are precisely 6 ways of splitting it: 1–12, 2–6, 3–4, 4–3, 6–2, and 12–1. We have seen the first three splits above. Write the expressions corresponding to the latter 3.

Exercise 19. What happens when you type `split [1 .. 12] : [5][2][8]?`

Exercise 20. Write a `split` expression to turn the sequence `[1 .. 120] : [120][8]` into a nested sequence with type `[3][4][10][8]`, keeping the elements in the same order. (**Hint:** Use nested comprehensions.)

2.5.6 Shifts and rotates

Common operations on sequences include shifting and rotating them. Cryptol supports both versions with left/right variants.

Exercise 21. Experiment with the following expressions:

```
[1, 2, 3, 4, 5] >> 2
[1, 2, 3, 4, 5] >> 10
[1, 2, 3, 4, 5] << 2
[1, 2, 3, 4, 5] << 10
[1, 2, 3, 4, 5] >>> 2
[1, 2, 3, 4, 5] >>> 10
[1, 2, 3, 4, 5] <<< 2
[1, 2, 3, 4, 5] <<< 10
```

Notice that shifting/rotating always returns a sequence precisely the same size as the original.

Exercise 22. Let xs be a sequence of length n . What is the result of rotating xs left or right by a multiple of n ?

2.6 Words revisited

In Section 2.3 we have introduced numbers as a distinct value type in Cryptol. In fact, a number in Cryptol is nothing but a finite sequence of bits, so words are not a separate type. For instance, the literal expression 42 is precisely the same as the bit-sequence corresponding to `[True, False, True, False, True, False]`.

Exercise 23. Explain why 42 is the same as `[True, False, True, False, True, False]`. Is Cryptol little-endian, or big-endian?

Exercise 24. Try out the following words: (**Hint:** It might help to use `:set base=2` to see the bit patterns.)

```
12
12 # [False]
[False, False] # 12
[True, False] # 12
12 # [False, True]
32
12 # 32
[True, False, True, False, True, False] == 42
```

Exercise 25. What is the type of 0? Use the `:t` command to find this out. (Type `:t 0` at the prompt.) Are there any other elements of this type? What are the elements of the type `[2]`?

Defaulting and explicit types Top level polymorphic constants in a Cryptol program are subject to *defaulting*, meaning that Cryptol will use the fewest number of bits necessary to represent them. Users can override this by giving an explicit type signature.

Exercise 26. Try the following expressions:

2.7. Characters and strings

```
:t 42
:t 42 : [9]
:t 42 : [3]
```

Can you jam more bits in a word than is potentially possible in Cryptol? Compare this behavior to a typical C expression: `(char) 9999`.

Exercise 27. Since words are sequences, the sequence functions from Exercise 2.5–15 apply to words as well. Try out the following examples and explain the outputs you observe:

```
take' {3} 0xFF
take' {3} (12:[6])
drop' {3} (12:[6])
split' {3} (12:[6])
groupBy' {3} (12:[6])
```

Recall that the notation `12:[6]` means the constant 12 with the type precisely 6-bits wide.

Exercise 28. Try Exercise 27, this time with the constant `12:[12]`. Do any of the results change? Why?

Shifts and rotates on words Consider what happens if we shift a word, say `12:[6]` by one to the right:

```
(12:[6]) >> 1
= [False, False, True, True, False, False] >> 1
= [False, False, False, True, True, False]
= 6
```

That is shifting-right by one effectively divides the word by 2. This is due to Cryptol’s “big-endian” representation of numbers⁴.

Exercise 29. Try the following examples of shifting/rotating words:

```
(12:[8]) >> 2
(12:[8]) << 2
```

Little-endian vs Big-endian The discussion of endianness comes up often in computer science, with no clear winner. Since Cryptol allows indexing from the beginning or the end of a (finite) sequence, you can access the 0th (least-significant) bit of a sequence k with $k!0$, the 1st bit with $k!1$, and so on.

2.7 Characters and strings

Strictly speaking Cryptol does *not* have characters and strings as a separate type. However, Cryptol does allow characters in programs, which simply correspond to their ASCII equivalents. Similarly, strings are merely sequences of characters, i.e., sequences of words. The following examples illustrate:

⁴This is a significant change from Cryptol version 1, which interpreted the leftmost element of a sequence as the lowest-ordered bit (and thus shifting right was multiplying by 2, and shifting left was dividing by 2). The way it is handled now matches the traditional interpretation.

```

Cryptol> :set base=10
Cryptol> :set ascii=off
Cryptol> 'A'
65
Cryptol> "ABC"
[65, 66, 67]
Cryptol> :set ascii=on
Cryptol> "ABC"
"ABC"
Cryptol> :set ascii=off
Cryptol> ['A' .. 'Z']
[65, 66, 67, 68, 69, 70, 71, 72, 73, 74, 75, 76, 77, 78, 79,
 80, 81, 82, 83, 84, 85, 86, 87, 88, 89, 90]
Cryptol> :set ascii=on
Cryptol> ['A' .. 'Z']
"ABCDEFGHIJKLMNOPQRSTUVWXYZ"

```

Note. This is the reason why we have to use the `:set ascii=on` command to print ASCII strings. Otherwise, Cryptol will not have enough information to tell numbers from characters.

Since characters are simply 8-bit words, you can do word operations on them; including arithmetic:

```

Cryptol> 'C' - 'A'
2

```

2.8 Records: Named collections

In Cryptol, records are simply collections of named fields. In this sense, they are very similar to tuples (Section 2.4), which can be thought of records without field names⁵. Like a tuple, the fields of a record can be of any type. We construct records by listing the fields inside curly-braces, separated by commas. We project fields out of a record with the usual dot-notation. Note that the order of fields in a record is immaterial.

Record equality is defined in the standard fashion. Two records are equal if they have the same number of fields, if their field names are identical, if identically named fields are of comparable types and have equal values.

Exercise 30. Type in the following expressions and observe the output:

```

{xCoord = 12:[32], yCoord = 21:[32]}
{xCoord = 12:[32], yCoord = 21:[32]}.yCoord
{name = "Cryptol", address = "Galois"}
{name = "Cryptol", address = "Galois"}.address
{name = "test", coords = {xCoord = 3:[32], yCoord = 5:[32]}}
{name = "test", coords = {xCoord = 3:[32], \
                          yCoord = 5:[32]}}.coords.yCoord
{x=True, y=False} == {y=False, x=True}

```

⁵In fact, the fields of a tuple *can* be accessed via the dot-notation, with their names being their 0-indexed position in the tuple. So $(1,2).1 == 2$.

2.9. The zero

You might find the command `:set ascii=on` useful in viewing the output.

Note. In larger Cryptol programs, records provide quite powerful abstraction mechanisms. In particular, record fields can contain polymorphic fields themselves, extracted and used at different types in the same expression. However, we will not need that level of functionality in our current study.

2.9 The zero

Before proceeding further, we have to take a detour and talk briefly about one of the most useful values in Cryptol: `zero`. The value `zero` inhabits every type in Cryptol, and stands for the value that consists of all `False` bits. The following examples should illustrate the idea:

```
Cryptol> zero : Bit
False
Cryptol> zero : [8]
0
Cryptol> zero : ([8], Bit)
(0, False)
Cryptol> zero : [8][3]
[0, 0, 0, 0, 0, 0, 0, 0]
Cryptol> zero : [3](Bit, [4])
[(False, 0), (False, 0), (False, 0)]
Cryptol> zero : {xCoord : [12], yCoord : [5]}
{xCoord=0, yCoord=0}
```

On the other extreme, the value `zero` combined with the complement operator `~` gives us values that consist of all all `True` bits:

```
Cryptol> ~zero : Bit
True
Cryptol> ~zero : [8]
255
Cryptol> ~zero : ([8], Bit)
(255, True)
Cryptol> ~zero : [8][3]
[7, 7, 7, 7, 7, 7, 7, 7]
Cryptol> ~zero : [3](Bit, [4])
[(True, 15), (True, 15), (True, 15)]
Cryptol> ~zero : {xCoord : [12], yCoord : [5]}
{xCoord=4095, yCoord=31}
```

Exercise 31. We said that `zero` inhabits all types in Cryptol. This also includes functions. What do you think the appropriate zero value for a function would be? Try out the following examples:

```
(zero : ([8] -> [3])) 5
(zero : Bit -> {xCoord : [12], yCoord : [5]}) True
```

2.10 Arithmetic

Cryptol supports the usual binary arithmetic operators $+$, $-$, $*$, $^{\wedge}$ (exponentiate), $/$ (integer division), $\%$ (integer modulus), along with *ceiling* logarithm base 2 $\text{lg}2$ and binary min and max .

The important thing to remember is that all arithmetic in Cryptol is modular, with respect to the underlying word size. As a consequence, there is no such thing as an overflow/underflow in Cryptol, as the result will be always guaranteed to fit in the resulting word size. While this is very handy for most applications of Cryptol, it requires some care if overflow has to be treated explicitly.

Exercise 32. What is the value of $1+1$? Surprised?

Exercise 33. What is the value of $1+(1:[8])$? Why?

Exercise 34. What is the value of $3 - 5$? How about $(3 - 5) : [8]$?

Note. Cryptol supports subtraction both as a binary operator, and as a unary operator. When used in a unary fashion (a.k.a., unary minus), it simply means subtraction from 0. For instance, -5 precisely means $0-5$, and is subject to the usual modular arithmetic rules.

Exercise 35. Try out the following expressions:

```
2 / 0
2 % 0
3 + (if 3 == 2+1 then 12 else 2/0)
3 + (if 3 != 2+1 then 12 else 2/0)
lg2 (-25)
```

In the last expression, remember that unary minus will be done in a modular fashion. What is the modulus used for this operation?

Exercise 36. Division truncates down. Try out the following expressions:

```
(6 / 3, 6 % 3)
(7 / 3, 7 % 3)
(8 / 3, 8 % 3)
(9 / 3, 9 % 3)
```

What is the relationship between $/$ and $\%$?

Exercise 37. What is the value of $\text{min } 5 (-2)$? Why? Why are the parentheses necessary?

Exercise 38. How about $\text{max } 5 (-2:[8])$? Why?

Exercise 39. Write an expression that computes the sum of two sequences $[1 .. 10]$ and $[10, 9 .. 1]$.

Comparison operators Cryptol supports the comparison operators $==$, $!=$, $>$, $>=$, $<$, $<=$, with their usual meanings.

Exercise 40. Try out the following expressions:

```
((2 >= 3) || (3 < 6)) && (4 == 5)
if 3 >= 2 then True else 1 < 12
```

2.10. Arithmetic

Enumerations, revisited In Exercise 2.5–13, we wrote the infinite enumeration starting at 1 using an explicit type as follows:

```
[(1:[32]) ... ]
```

As expected, Cryptol evaluates this expression to:

```
[1, 2, 3, 4, 5, ...]
```

However, while the output suggests that the numbers are increasing all the time, that is just an illusion! Since the elements of this sequence are 32-bit words, eventually they will wrap over, and go back to 0. (In fact, this will happen precisely at the element $2^{32} - 1$, starting the count at 0 as usual.) We can observe this much more simply, by using a smaller bit size for the constant 1:

```
Cryptol> [(1:[2])...]
[1, 2, 3, 0, 1 ...]
Cryptol> take' {20} [(1:[2])...]
[1, 2, 3, 0, 1, 2, 3, 0, 1, 2, 3, 0, 1, 2, 3, 0, 1, 2, 3, 0]
```

We still get an infinite sequence, but the numbers will repeat themselves eventually. Note that this is a direct consequence of Cryptol's modular arithmetic.

There is one more case to look at. What happens if we completely leave out the signature?

```
Cryptol> [1 ... ]
[1, 0, 1, 0, 1, ...]
```

In this case, Cryptol figured out that the number 1 requires precisely 1-bits, and hence the arithmetic is done modulo $2^1 = 2$, giving us the sequence 1-0-1-0 In particular, an enumeration of the form:

```
[k ..]
```

will be treated as if the user has written:

```
[k, (k+1) ..]
```

and type inference will assign the smallest bit-size possible to represent k .

Note. if the user evaluates the value of $k+1$, then the result may be different. For example, `[1, 1+1 ...]` results in the `[1, 0, 1 ...]` behavior, but `[1, 2 ...]` adds another bit, resulting in `[1, 2, 3, 0, 1, 2, 3 ...]`. If Cryptol evaluates the value of $k+1$, the answer is modulo k , so another bit is not added. For the curious, this subtle behavior was introduced to allow the sequence of all zeros to be written `[0 ...]`.

Exercise 41. Remember from Exercise 2.6–25 that the constant 0 requires 0-bits to represent. Based on this, what is the value of the enumeration `[0 ...]`? What about `[0 ...]`? Surprised?

Exercise 42. What is the value of `[1 .. 10]`? Explain in terms of the above discussion on modular arithmetic.

2.11 Types

Cryptol’s type system is one of its key features⁷. You have seen that types can be used to specify the exact width of values, or shapes of sequences using a rich yet concise notation. In some cases, it may make sense to omit a type signature and let Cryptol *infer* the type for you. At the interpreter, you can check what type Cryptol inferred with the `:t` command.

2.11.1 Monomorphic types

A monomorphic type is one that represents a concrete value. Most of the examples we have seen so far falls into this category. Below, we review the basic Cryptol types that make up all the monomorphic values in Cryptol.

Bits There are precisely two bit values in Cryptol: `True` and `False`. The type itself is written `Bit`. When we want to be explicit, we can write it as follows: `(2 >= 3) : Bit`. However, with type inference writing the `Bit` type explicitly is almost never needed.

Words A word type is written `[n]`, where n is a fixed non-negative constant. The constant can be as large (or small) as you like. So, you can talk about 2-bit quantities `[2]`, as well as 384-bit ones `[384]`, or even odd sizes like 17 `[17]`, depending on the needs of your application. When we want to be explicit about the type of a value, we say `5 : [8]`. If we do not specify a size, Cryptol’s type inference engine will pick the appropriate value depending on the context. Recall from Section 2.6 that a word is, in fact, a sequence of bits. Hence, an equivalent (but verbose) way to write the type `[17]` is `[17]Bit`, which we would say in English as “a sequence of length 17, whose elements are Bits.”

Tuples A tuple is a heterogeneous collection of arbitrary number of elements. Just like we write a tuple value by enclosing it in parentheses, we write the tuple type by enclosing the component types in parentheses, separated by commas: `(3, 5, True) : ([8], [32], Bit)`. Tuples’ types follow the same structure: `(2, (False, 3), 5) : ([8], (Bit, [32]), [32])`. A tuple component can be any type: a word, another tuple, sequence, record, etc. Again, type inference makes writing tuple types hardly ever necessary.

Sequences A sequence is simply a collection of homogeneous elements. If the element type is τ , then we write the type of a sequence of n elements as: `[n] τ` . Note that τ itself can be a sequence itself. For instance, the type: `[12][3][6]` reads as follows: A sequence of 12 elements, each of which is a sequence of 3 elements, each of which is a 6-bit wide word.

The type of an infinite sequence is written `[inf] τ` , where τ is the type of the elements.

Exercise 43. What is the total number of bits in the type `[12][3][6]`?

Exercise 44. How would you write the type of an infinite sequence where each element itself is an infinite sequence of 32 bit words? What is the total bit size of this type?

⁷The Cryptol type system is based on the traditional Hindley-Milner style, extended with size types and arithmetic predicates [6, 7, 8]

2.11. Types

Records A record is a heterogeneous collection of arbitrary number of labeled elements. In a sense, they generalize tuples by allowing the programmer to give explicit names to fields. The type of a record is written by enclosing it in braces, separated by commas: `{x : [32], y : [32]}`. Records can be nested and can contain arbitrary types of elements (records, sequences, functions, etc.).

2.11.2 Polymorphic types

Our focus so far has been on monomorphic types—the types that concrete Cryptol values (such as `True`, `3`, or `[1, 2]`) can have. If all we had were monomorphic types, however, Cryptol would be a very verbose and boring language. Instead, we would like to be able to talk about collections of values, values whose types are instances of a given polymorphic type. This facility is especially important when we define functions, a topic we will get to shortly. In the mean time, we will look at some of the polymorphic primitive functions Cryptol provides to get a feeling for Cryptol’s polymorphic type system.

The tale of `tail` Cryptol’s built in function `tail` allows us to drop the first element from a sequence, returning the remainder:

```
Cryptol> tail [1 .. 5]
[2, 3, 4, 5]
Cryptol> tail [(False, (1:[8])), (True, 12), (False, 3)]
[(True, 12), (False, 3)]
Cryptol> tail [ (1:[16])... ]
[2, 3, 4, 5, 6, ...]
```

What exactly is the type of `tail`? If we look at the first example, one can deduce that `tail` must have the type:

```
tail : [5][8] -> [4][8]
```

That is, it takes a sequence of length 5, containing 8-bit values, and returns a sequence that has length 4, containing 8-bit values. (The type `a -> b` denotes a function that takes a value of type `a` and delivers a value of type `b`.)

However, the other example uses of `tail` above suggest that it must have the following types, respectively:

```
tail : [10][32] -> [9][32]
tail : [3](Bit, [8]) -> [2](Bit, [8])
tail : [inf][16] -> [inf][16]
```

As we have emphasized before, Cryptol is strongly-typed, meaning that every entity (whether a Cryptol primitive or a user-defined function) must have a well-defined type. Clearly, the types we provided for `tail` above are quite different from each other. In particular, the first example uses numbers as the element type, while the second has tuples. So, how can `tail` be assigned a type that will make it work on all these inputs?

If you are familiar C++ templates or Java generics, you might think that Cryptol has some sort of an overloading mechanism that allows one to define functions that can work on multiple types. While templates and generics do provide a mental model, the correspondence is not very strong. In particular, we never write multiple definitions for the same function in Cryptol, i.e., there is no ad-hoc overloading. However, what Cryptol has is a much stronger notion: polymorphism, as would be advocated by languages such as Haskell or ML [12, 14].

Here is the type of `tail` in Cryptol:

```
Cryptol> :t tail
tail : {a, b} [a+1]b -> [a]b
```

This is quite a different type from what we have seen so far. In particular, it is a polymorphic type, one that can work over multiple concrete instantiations of it. Here's how we read this type in Cryptol:

tail is a polymorphic function, parameterized over *a* and *b*. The input is a sequence that contains *a*+1 elements. The elements can be of an arbitrary type *b*, there is no restriction on their structure. The result is a sequence that contains *a* elements, where the elements themselves has the same type as those of the input.

In the case for `tail`, the parameter *a* is a size-parameter (since it describes the size of a sequence), while *b* is a shape-parameter, since it describes the shape of elements. The important thing to remember is that each use of `tail` must instantiate the parameters *a* and *b* appropriately. Let's see how the instantiations work for our running examples:

<code>[a+1]b -> [a][b]</code>	<code>a</code>	<code>b</code>	Notes
<code>[5][8] -> [4][8]</code>	4	<code>[8]</code>	$a+1 = 5 \Rightarrow a = 4$
<code>[10][32] -> [9][32]</code>	9	<code>[32]</code>	$a+1 = 10 \Rightarrow a = 9$
<code>[3](Bit, [8]) -> [2](Bit, [8])</code>	2	<code>(Bit, [8])</code>	The type <i>b</i> is now a tuple
<code>[inf][16] -> [inf][16]</code>	<code>inf</code>	<code>[16]</code>	$a+1 = \text{inf} \Rightarrow a = \text{inf}$

In the last instantiation, Cryptol knows that $\infty - 1 = \infty$, allowing us to apply `tail` on both finite and infinite sequences. The crucial point is that an instantiation must be found that satisfies the required match. It is informative to see what happens if we apply `tail` to an argument where an appropriate instantiation can not be found:

```
Cryptol> tail True
[error] at <interactive>:1:1--1:10:
Type mismatch:
  Expected type: Bit
  Inferred type: [1 + ?a]?b
```

Cryptol is telling us that it cannot match the types `Bit` and the sequence `[a+1]b`, causing a type error statically at compile time. (The funny notation of `?a` and `?b` are due to how type instantiations proceed under-the-hood. While they look funny at first, you soon get used to the notation.)

We should emphasize that Cryptol polymorphism uniformly applies to user-defined functions as well, as we shall see in Section 2.12.

Exercise 45. Use the `:t` command to find the type of `groupBy`. For each use case below, find out what the instantiations of its type variables are, and justify why the instantiation works. Can you find an instantiation in all these cases?

```
groupBy' {3} [1..9]
groupBy' {3} [1..12]
groupBy' {3} [1..10] : [3][2][8]
groupBy' {3} [1..10]
```

Is there any way to make the last example work by giving a type signature?

2.11. Types

2.11.3 Predicates

In the previous section we have seen how polymorphism is a powerful tool in structuring programs. Cryptol takes the idea of polymorphism on sizes one step further by allowing predicates on sizes. To illustrate the notion, consider the type of the Cryptol primitive `take`:

```
Cryptol> :t take
take : {front, back, elem} (fin front) => [front + back]elem
      -> [front]elem
```

The type of `take` says that it is parameterized over `front` and `back`, `front` must be a finite value, it takes a sequence `front + back` long, and returns a sequence `front` long.

The impact of this predicate shows up when we try to take more than what is available:

```
Cryptol> take' {10} [1..5]
[error] at <interactive>:1:1--1:17:
  Unsolved constraint:
    0 == 5 + ?a
    arising from matching types at <interactive>:1:1--1:17
```

Cryptol is telling us that it is unable to satisfy this instantiation (since `front` is 10 and the sequence has 5 elements).

In general, type predicates exclusively describe arithmetic constraints on type variables. Cryptol does not have a general-purpose dependent type system, but a *size-polymorphic type system*. Often type variables' values are of finite size, indicated with the constraint `fin a`, otherwise no constraint is mentioned or an explicit `inf a` is denoted, and the variables' values are unbounded. Arithmetic relations are arbitrary relations over all type variables, such as $2*a+b \geq c$. We shall see more examples as we work through programs later on.

Exercise 46. Write a predicate that allows a word of size 128, 192, or 256, but nothing else.

Note. Type inference in the presence of arithmetic predicates is an undecidable problem [7]. This implies that there is no algorithm to decide whether a given type is inhabited, amongst other things. In practical terms, we might end up writing programs with arbitrarily complicated predicates (e.g., this “type contains all solutions to Fermat’s last equation” or “this type contains all primes between two large numbers”) without Cryptol being able to simplify them, or notice that there is no instantiation that will ever work. Here is a simple example of such a type:

```
{k} (2 >= k, k >= 5) => [k]
```

While a moment of pondering suffices to conclude that there is no such value in this particular case, there is no algorithm to decide this problem in general.

That being said, Cryptol’s type inference and type checking algorithms are well-tuned to the use cases witnessed in the types necessary for cryptographic algorithms. Moreover, Cryptol uses a powerful SMT solver capable of reasoning about complex arithmetic theories within these algorithms.

2.11.4 Why typed?

There is a spectrum of type systems employed by programming languages, all the way from completely untyped to fancier dependently typed languages. There is no simple answer to the question what type system

is the best? It depends on the application domain. We have found that Cryptol’s size-polymorphic type system is a good fit for programming problems that arise in the domain of cryptography. The bit-precise type system makes sure that we never pass an argument that is 32-bits wide in a buffer that can only fit 16. The motto is: *Well typed programs do not go wrong*.

In practical terms, this means that the type system catches most of the common mistakes that programmers tend to make. Size-polymorphism is a good fit for Cryptol, as it keeps track of the most important invariant in our application domain: making sure the sizes of data can be very precisely specified and the programs can be statically guaranteed to respect these sizes.

Opponents of type systems typically argue that the type system gets in the way⁸. It is true that the type system will reject some programs that makes perfect sense. But what is more important is that the type system will reject programs that will indeed go wrong at run-time. And the price you pay to make sure your program type-checks is negligible, and the savings due to type-checking can be enormous.

The crucial question is not whether we want type systems, but rather what type system is the best for a given particular application domain. We have found that the size-polymorphic type system of Cryptol provides the right balance for the domain of cryptography and bit-precise programming problems it has been designed for [11].

2.12 Defining functions

So far, we used Cryptol as a calculator: we typed in expressions and it gave us answers. This is great for experimenting, and exploring Cryptol itself. The next fundamental Cryptol idiom is the notion of a function. You have already used built-in functions `+`, `take`, etc. Of course, users can define their own functions as well. Currently the Cryptol interpreter does not support defining functions, so you must define them in a file and load it, as in the next exercises.

Note. Reviewing the contents of Section 1.2 might help at this point. Especially the commands that will let you load files (`:l` and `:r`) in Cryptol.

Exercise 47. Type the following definition into a file and save it. Then load it in Cryptol and experiment.

```
increment : [8] -> [8]
increment x = x+1
```

In particular, try the following invocations:

```
increment 3
increment 255
increment 912
```

Do you expect the last call to type check?

Note. Note that we do not have to parenthesize the argument to `increment`, as in `increment(3)`. Function application is simply juxtaposition in Cryptol. However, you can put the parentheses if you want to, and you must use parentheses if you want to pass a negative argument (e.g., `increment(-2)` (recall 37)

⁸Another complaint is that “strong types are for weak minds.” We do not disagree here: Cryptol programmers want to use the type system so we do not have to think as hard about writing correct code as we would without strong types.

2.12. Defining functions

2.12.1 Local names: where-clauses

You can create local bindings in a `where` clause, to increase readability and give names to common subexpressions.

Exercise 48. Define and experiment with the following function:

```
twoPlusXY : ([8], [8]) -> [8]
twoPlusXY (x, y) = 2 + xy
  where xy = x * y
```

What is the signature of the function telling us?

Note. When calling `twoPlusXY`, you do need to parenthesize the arguments. But this is because you are passing it a tuple! The parentheses there are not for the application but rather in construction of the argument tuple. Cryptol does not automatically convert between tuples and curried application like in some other programming languages (e.g., one cannot pass a pair of type (a, b) to a function with type $a \rightarrow b \rightarrow c$).

Exercise 49. Comment out the type signature of `twoPlusXY` defined above, and let Cryptol infer its type. What is the inferred type? Why?

Exercise 50. Define a function with the following signature:

```
minMax4 : {a} (Cmp a) => [4]a -> (a, a)
```

such that the first element of the result is the minimum and the second element is the maximum of the given four elements. What happens when you try:

```
minMax4 [1 .. 4]
minMax4 [1 .. 5]
```

Why do we need the `(Cmp a)` constraint?

Exercise 51. Define a function `butLast` that returns all the elements of a given non-empty sequence, except for the last. How can you make sure that `butLast` can never be called with an empty sequence? (**Hint:** You might find the Cryptol primitive functions `reverse` and `width` useful.)

2.12.2 λ -expressions

One particular use case of a `where`-clause is to introduce a helper function. If the function is simple enough, though, it may not be worth giving it a name. A λ -expression fits the bill in these cases, where you can introduce an unnamed function as an expression. The syntax differs from ordinary definitions in two minor details: instead of the name we use the backslash or “wack” character, `\`, and the equals sign is replaced by an arrow `->`. (Since these functions do not have explicit names, they are sometimes referred to as “anonymous functions” as well. We prefer the term λ -expression, following the usual functional programming terminology [14].)

Below is an example of a λ -expression, allowing us to write functions inline:

```
Cryptol> f 8 where f x = x+1
9
Cryptol> (\x -> x+1) 8
9
```

λ -expressions are especially handy when you need to write small functions at the command line while interacting with the interpreter.

2.12.3 Using `zero` in functions

The constant `zero` comes in very handy in Cryptol whenever we need a polymorphic shape that consists of all `False` bits. The following two exercises utilize `zero` to define the functions `all` and `any` which, later in this book, you will find are very helpful functions for producing boolean values from a sequence.

Exercise 52. Write a function `all` with the following signature:

```
all : {n, a} (fin n) => (a -> Bit) -> [n]a -> Bit
```

such that `all f xs` returns `True` if all the elements in the sequence `xs` yield `True` for the function `f`. (**Hint:** Use a complemented `zero`.) You should see:

```
Cryptol> all eqTen [10, 10, 10, 10] where eqTen x = x == 10
True
Cryptol> all eqTen [10, 10, 10, 5] where eqTen x = x == 10
False
```

(The `where`-clause introduces a local definition that is in scope in the current expression. We will see the details in Section 2.12.) What is the value of `all f []` for any `f`? Is this reasonable?

Exercise 53. Write a function `any` with the following signature:

```
any : {n, a} (fin n) => (a -> Bit) -> [n]a -> Bit
```

such that `any f xs` returns `True` if any the elements in the sequence `xs` yield `True` for the function `f`. What is the value of `any f []`? Is this reasonable?

2.13 Recursion and recurrences

Cryptol allows both recursive function and value definitions. A recursive function is one that calls itself in its definition. Cryptol also allows the more general form of mutual recursion, where multiple functions can call each other in a cyclic fashion.

Exercise 54. Define two functions `isOdd` and `isEven` that each take a finite arbitrary sized word and returns a `Bit`. The functions should be mutually recursive. What extra predicates do you have to put on the `size`?

Exercise 55. While defining `isOdd` and `isEven` mutually recursively demonstrates the concept of recursion, it is not the best way of coding these two functions in Cryptol. Can you implement them using a constant time operation? (**Hint:** What is the least significant bit of an even number? How about an odd one?)

Recurrences While Cryptol does support recursion, the explicit recursive function style is typically discouraged: Arbitrary recursion is hard to compile down to hardware. A much better notion is that of *recurrences*. A recurrence is a way of cyclically defining a value, typically a stream. It turns out that most recursive functions can be written in a recurrence style as well, something that might first come as a surprise.

2.13. Recursion and recurrences

In particular, most recursive definitions arise from recurrence equations in cryptographic and data flow style programs, and Cryptol's comprehensions can succinctly represent these computations.

Exercise 56. In this exercise, we will define a function to compute the maximum element in a given sequence of numbers. Define the following function and load it into Cryptol:

```
maxSeq xs = ys ! 0
  where ys = [0] # [ max x y | x <- xs
                    | y <- ys
                    ]
```

What is the type of `maxSeq`? Try out the following calls:

```
maxSeq []
maxSeq [1 .. 10]
maxSeq ([10 .. 1] # [1 .. 10])
```

Patterns of recurrence The definition pattern for `ys` in the definition of `maxSeq` above is very common in Cryptol, and it is well worth understanding it clearly. The basic idea is to create a sequence of running results, for each prefix of the input.

Exercise 57. Define a variant of `maxSeq` (let us call it `maxSeq'`) which returns the sequence `ys`, instead of its last element.

Running results It is very instructive to look at the results returned by `maxSeq'` that you have just defined in Exercise 57:

```
Cryptol> maxSeq' []
[0]
Cryptol> maxSeq' [1 .. 10]
[0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10]
Cryptol> maxSeq' [10 .. 1]
[0, 10, 10, 10, 10, 10, 10, 10, 10, 10, 10]
Cryptol> maxSeq' [1, 3, 2, 4, 3, 5, 4, 7, 6, 8, 5]
[0, 1, 3, 3, 4, 4, 5, 5, 7, 7, 8, 8]
```

We clearly see the running results as they accumulate in `ys`. For the empty sequence, it only has `[0]` in it. For the monotonically increasing sequence `[1 .. 10]`, the maximum value keeps changing at each point, as each new element of `xs` is larger than the previous running result. When we try the sequence that always goes down (`[10, 9 .. 1]`), we find that the running maximum never changes after the first. The mixed input in the last call clearly demonstrates how the execution proceeds, the running maximum changing depending on the next element of `xs` and the running maximum so far. In the `maxSeq` function of Exercise 56, we simply project out the last element of this sequence, obtaining the maximum of all elements in the given sequence.

Folds The pattern of recurrence employed in `maxSeq` is an instance of what is known as a *fold* [1]. Expressed in Cryptol terms, it looks as follows:

```
ys = [i] # [ f (x, y) | x <- xs
              | y <- ys
              ]
```

where xs is some input sequence, i is the result corresponding to the empty sequence, and f is a transformer to compute the next element, using the previous result and the next input. This pattern can be viewed as generating a sequence of running values, accumulating them in ys . To illustrate, if xs is a sequence containing the elements $[x_1, x_2, x_3 \dots x_n]$, then successive elements of ys will be:

$$\begin{aligned} y_0 &= i \\ y_1 &= f(x_1, i) \\ y_2 &= f(x_2, y_1) \\ y_3 &= f(x_3, y_2) \\ &\dots \\ y_n &= f(x_n, y_{n-1}) \end{aligned}$$

Note how each new element of ys is computed by using the previous element and the next element of the input. The value i provides the seed. Consequently, ys will have one more element than xs does.

While loops An important use case of the above pattern is when we are interested in the final value of the accumulating values, as in the definition of `maxSeq`. When used in this fashion, the execution is reminiscent of a simple while loop that you might be familiar from other languages, such as C:

```
-- C-like Pseudo-code!
y = i;           // running value, start with i
idx = 0;        // walk through the xs "array" using idx
while(idx <= length (xs)) {
  y = f (xs[idx], y); // compute next elt using the previous
  ++idx;
}
return y;
```

Note. If the while-loop analogy does not help you, feel free to ignore it. It is not essential. The moral of the story is this: if you feel like you need to write a while-loop in Cryptol to compute a value dependent upon the values in a datatype, you probably want to use a fold-like recurrence instead.

Exercise 58. Define a function that sums up a given sequence of elements. The type of the function should be:

```
sumAll : {n, a} (fin n, fin a) => [n][a] -> [a]
```

(**Hint:** Use the folding pattern to create a sequence containing the partial running sums. What is the last element of this sequence?) Try it out on the following examples:

```
sumAll []
sumAll [1]
sumAll [1, 2]
sumAll [1, 2, 3]
sumAll [1, 2, 3, 4]
sumAll [1, 2, 3, 4, 5]
sumAll [1 .. 100]
```

2.14. Stream equations

Be mindful that if you do not specify the width of the result that you may get unexpected answers.

Exercise 59. Define a function `elem` with the following signature:

```
elem : {n, t} (fin n, Cmp t) => (t, [n]t) -> Bit
```

such that `elem (x, xs)` returns `True` if `x` appears in `xs`, and `False` otherwise.

Generalized folds The use of `fold` we have seen above is the simplest use case for recurrences in Cryptol. It is very common to see Cryptol programs employing some variant of this idea for most of their computations.

Exercise 60. Define the sequence of Fibonacci numbers `fib`s, so that `fib`s @ `n` is the n^{th} Fibonacci number [20]. You can assume 32 bits is sufficient for representing these fast growing numbers. (**Hint:** Use a recurrence where the seed consists of two elements.)

2.14 Stream equations

Most cryptographic algorithms are described in terms of operations on bit-streams. A common way of depicting operations on bit-streams is using a *stream equation*, as shown in Figure 2.1:




Figure 2.1: Equation for producing a stream of `as`

In this diagram the stream is seeded with four initial values (`3F`, `E2`, `65`, `CA`). The subsequent elements (`new`) are appended to the stream, and are computed by xor-ing the current stream element with two additional elements extracted from further into the stream. The output from the stream is a sequence of values, known as 'a's'.

The Cryptol code corresponding to this stream equation is:

```
as = [0x3F, 0xE2, 0x65, 0xCA] # new
where
  new = [ a ^ b ^ c | a <- as
           | b <- drop' {1} as
           | c <- drop' {3} as ]
```

Exercise 61. Write the Cryptol code corresponding to the stream equation in Figure 2.2:




Figure 2.2: Equation for producing a stream of as from an initial seed and an input stream.

2.15 Type synonyms

Types in Cryptol can become fairly complicated, especially in the presence of records. Even for simple types, meaningful names should be used for readability and documentation. Type synonyms allow users to give names to arbitrary types. In this sense, they are akin to `typedef` declarations in C [9]. However, Cryptol's type synonyms are significantly more powerful than C's `typedef`'s, since they can be parameterized by other types, much like in Haskell [14].

Here are some simple type synonym definitions:

```

type Word8      = [8]
type CheckedWord = (Word8, Bit)
type Point a    = {x : [a], y : [a]}
  
```

Type synonyms are either unparameterized (as in `word8` and `CheckedWord`, or parameterized with other types (as in `Point`). Synonyms may depend upon other synonyms, as in the `CheckedWord` example. Once the synonym is given, it acts as an additional name for the underlying type, making it much easier to read and maintain.

For instance, we can write the function that returns the x-coordinate of a point as follows:

```

xCoord : {a} Point a -> [a]
xCoord p = p.x
  
```

Note that type synonyms, while maintained within the type and value context shown via the `:browse` command, are *value-based*, not *name-based*. When viewed from the types-as-sets interpretation, two types in Cryptol are synonymous if their values happen to be equal.

For example, consider the following declarations:

```

type Word8      = [8]
type Word8'     = [8]
type B          = Word8
type A          = B
type WordPair   = (Word8, Word8')
type WordPair' = (Word8', Word8)
  
```

```

foo: Word8 -> Bit
foo x = True
  
```

```

bar: Word8' -> Bit
bar x = foo x
  
```


2.16. Type classes

Within this type context, while six *names* are declared, only *two* types are declared (`[8]` and the pair `([8], [8])`). Likewise, the function types of `foo` and `bar` are identical, thus `bar` can call `foo`.

Exercise 62. Define a type synonym for 3-dimensional points and write a function to determine if the point lies on any of the 3 axes.

Predefined type synonyms The following type synonyms are predefined in Cryptol:

```
type Bool = Bit
type Char = [8]
type String n = [n]Char
type Word n = [n]
```

For instance, a `String n` is simply a sequence of precisely `n` 8-bit words.

2.16 Type classes

Type classes are a way of describing behaviors shared by multiple types. As an example, consider the type of the function `==`:

```
Cryptol> :t (==)
== : {a} (Cmp a) => a -> a -> Bit
```

This operator type is interpreted “equality is an operator that takes two objects of any single type that can be compared and returns a `Bit`.”

Cryptol defines exactly two basic type classes: `Cmp` and `Arith`. These appear in the type signature of operators and functions that require them. If a function you define calls, for example, `+`, on two arguments both of type `a`, the type constraints for `a` will include `(Arith a)`.

The `Cmp` typeclass includes the binary relation operators `<`, `>`, `<=`, `>=`, `==`, and `!=`, as well as the binary functions `min` and `max`. Note that equality is defined on function types (i.e., `a b (Cmp b) => (a -> b) -> (a -> b) -> a -> Bit`). Unlike in many other languages, equality and comparison are bundled into a single typeclass.

The `Arith` typeclass include the binary operators `+`, `-`, `*`, `/`, `%`, `^^`, as well as the unary operators `lg2` and `-`.

Exercise 63. Without including an explicit type declaration, define a function that Cryptol infers has the following type:

```
cmpArith : {a,b} (Cmp a, Arith b) => a -> a -> b -> b
```

2.17 Type vs. value variables

Its powerful type system is one of the key features of Cryptol. We have encountered many aspects of types already. You may have noticed, in functions such as `groupBy`, that when you call a function in Cryptol, there are two kinds of parameters you can pass: *value variables* and *type variables*.

Consider the `groupBy` function that we previously examined in 45. Recall that `groupBy`'s type is:

```
groupBy : {each, parts, elem} (fin each) =>
  [parts * each]elem -> [parts][each]elem
```

When applying `groupBy`, one typically specifies a concrete value for the formal parameter `parts`:

```
Cryptol> groupBy' {parts=3} [1..12]
[[1, 2, 3, 4], [5, 6, 7, 8], [9, 10, 11, 12]]
```

In this example, the term `' {parts=3}` passes 3 to the `parts` type variable argument, and the `[1..12]` is passing a sequence as the first (and only) *value argument*, `elem`.

A *value variable* is the kind of variable you are used to from normal programming languages. These kinds of variable represent a normal run-time value.

A *type variable*, on the other hand, allows you to express interesting (arithmetic) constraints on *types*. These variables express things like lengths of sequences or relationships between lengths of sequences. Type variable values are computed statically—they never change at runtime⁹.

2.17.1 Positional vs. named type arguments

Cryptol permits type variables to be passed either by name (as in `' {parts=3}` above), or by position (leaving out the name). For functions you define, the position is the order in which the type variables are declared in your function's type signature. If you are not sure what that is, you can always use the `:t` command to find out the position of type variables.

For example:

```
Cryptol> :t groupBy
groupBy : {each, parts, elem}
  (fin each) => [parts * each]elem
  -> [parts][each]elem
```

tells us that that `parts` is in the second position of `groupBy`'s type signature, so the positional-style call equivalent to our example is:

```
Cryptol> groupBy' {_,3} [1..12]
```

Note the use of an underscore in order to pass 3 in the second position. Positional arguments are most often used when the type argument is the first argument and when the name of the argument does not add clarity. The `groupBy' {_,3}` is not as self-explanatory as `groupBy' {parts=3}`. On the other hand, our use of positional arguments to `take` in previous chapters is very clear, as in:

```
Cryptol> take' {3} [1..12]
[1, 2, 3]
```

Tip. Cryptol programs that use named arguments are more maintainable and robust during program evolution. E.g., you can reorder parameters or refactor function definitions much more easily if invocations of those functions use named, rather than positional, arguments.

⁹In this way, they are similar (but more expressive than) templates in languages like C++ or Java. If you want to learn more about this area, look up the term “type-level naturals”.

2.17.2 Type context vs. variable context

You have seen, in the discussion of type variables above, that Cryptol has two kinds of variables – type variables and value variables. Type variables normally show up in type signatures, and value variables normally show up in function definitions. Sometimes you may want to use a type variable in a context where value variables would normally be used. To do this, use the backtick character `.

The definition of the built-in `width` function is a good example of the use of backtick:

```
width : {bits,len,elem} (fin len, fin bits, bits >= width len) =>
      [len] elem -> [bits]
width _ = `len
```

Tip. Note there are some subtle things going on in the above definition of `width`. First, arithmetic constraints on types are position-independent; properties of formal parameters early in a signature can depend upon those late in a signature. Second, type constraints can refer to not only other functions, but recursively to the function that is being defined (either directly, or transitively).

Type constraints can get pretty crazy in practice, especially deep in the signature of crypto code subsystems. Our suggestion is that you should not chase the dragon’s tail of feedback from the typechecker in attempting to massage your specification’s types for verification. Instead, think carefully about the meaning and purpose of the concepts in your specification, introduce appropriate type synonyms, and ensure that the specification is clear and precise. Trust that the interpreter and the verifier will do the right thing.

The bounds in a finite sequence literal (such as `[1 .. 10]`) in Cryptol are type-level values because the length of a sequence is part of its type. Only type-level values can appear in a finite sequence definition. You cannot write `[a .. b]` where either `a` or `b` are arguments to a function. On the other hand, an infinite sequence’s type is fixed (`(!inf]a)`), so the initial value in an infinite sequence can be a runtime variable or a type variable, but type variables are escaped here using a `.

This is probably obvious, but there is no way to get a value variable to appear in a type context. Types must be known at “compile time,” and (non-literal) values are not, so there is no way to use them in that way.

2.17.3 Inline argument type declarations

So far when we have defined a function, we have declared the type of its arguments and its return value in a separate type declaration. When you are initially writing code, you might not know exactly what a function’s full type is (including the constraints), but you may know (and need to express) the types of the function’s arguments. Cryptol’s syntax for this should look familiar:

```
addBytes (x:[8]) (y:[8]) = x + y
```

This defines a function that takes two bytes as input, and returns their sum. Note that the use of `()`’s is mandatory.

Here is a more interesting example:

```
myWidth (x:[w]a) = `w
```

2.18 The road ahead

In this introductory chapter, we have seen essentially all of the language elements in Cryptol. The concepts go deeper, of course, but you now have enough knowledge to tackle large Cryptol programming tasks. As

with any new language, the more exercises you do, the more you will feel comfortable with the concepts. In fact, we will take that point of view in the remainder of this document to walk you through a number of different examples (both small and large), employing the concepts we have seen thus far.

Chapter 3

Classic ciphers

Modern cryptography has come a long way. In his excellent book on cryptography, Singh traces it back to at least 5th century B.C., to the times of Herodotus and the ancient Greeks [15]. That's some 2500 years ago, and surely we do not use those methods anymore in modern day cryptography. However, the basic techniques are still relevant for appreciating the art of secret writing.

Shift ciphers construct the ciphertext from the plaintext by means of a predefined *shifting* operation, where the cipherkey of a particular shift algorithm defines the shift amount of the cipher. Transposition ciphers work by keeping the plaintext the same, but *rearrange* the order of the characters according to a certain rule. The cipherkey is essentially the description of how this transposition is done. Substitution ciphers generalize shifts and transpositions, allowing one to substitute arbitrary codes for plaintext elements. In this chapter, we will study several examples of these techniques and see how we can code them in Cryptol.

In general, ciphers boil down to pairs of functions *encrypt* and *decrypt* which “fit together” in the appropriate way. Arguing that a cryptographic function is *correct* is subtle.

Correctness of cryptography is determined by cryptanalyses by expert cryptographers. Each kind of cryptographic primitive (i.e., a hash, a symmetric cipher, an asymmetric cipher, etc.) has a set of expected properties, many of which can only be discovered and proven by hand through a lot of hard work. Thus, to check the correctness of a cryptographic function, a best practice for Cryptol use is to encode as many of these properties as one can in Cryptol itself and use Cryptol's validation and verification capabilities, discussed later in chapter 5. For example, the fundamental property of most ciphers is that encryption and decryption are inverses of each other.

To check the correctness of an *implementation* I of a cryptographic function C means that one must show that the implementation I behaves as the specification (C) stipulates. In the context of cryptography, the minimal conformance necessary is that I 's output *exactly* conforms to the output characterized by C . But just because a cryptographic implementation is *functionally correct* does not mean it is *secure*. The subtleties of an implementation can leak all kinds of information that harm the security of the cryptography, including abstraction leaking of sensitive values, timing attacks, side-channel attacks, etc. These kinds of properties cannot currently be expressed or reasoned about in Cryptol.

Also, Cryptol does *not* give the user any feedback on the *strength* of a given (cryptographic) algorithm. While this is an interesting and useful feature, it is not part of Cryptol's current capabilities.

3.1 Caesar's cipher

Caesar's cipher (a.k.a. Caesar's shift) is one of the simplest ciphers. The letters in the plaintext are shifted by a fixed number of elements down the alphabet. For instance, if the shift is 2, A becomes C, B becomes D, and so on. Once we run out of letters, we circle back to A; so Y becomes A, and Z becomes B. Coding Caesar's cipher in Cryptol is quite straightforward (recall from Section 2.15 that a `String n` is simply a sequence of `n` 8-bit words.):

```
caesar : {n} ([8], String n) -> String n
caesar (s, msg) = [ shift x | x <- msg ]
  where map      = ['A' .. 'Z'] <<< s
        shift c = map @ (c - 'A')
```

In this definition, we simply get a message `msg` of type `String n`, and perform a `shift` operation on each one of the elements. The `shift` function is defined locally in the `where`-clause. To compute the shift, we first find the distance of the letter from the character 'A' (via `c - 'A'`), and look it up in the mapping imposed by the shift. The `map` is simply the alphabet rotated to the left by the shift amount, `s`. Note how we use the enumeration `['A' .. 'Z']` to get all the letters in the alphabet.

Exercise 1. What is the map corresponding to a shift of 2? Use Cryptol's `<<<` to compute it. You can use the command `:set ascii=on` to print strings in ASCII, like this:

```
Cryptol> :set ascii=on
Cryptol> "Hello World"
"Hello World"
```

Why do we use a left-rotate, instead of a right-rotate?

Exercise 2. Use the above definition to encrypt the message "ATTACKATDAWN" by shifts 0, 3, 12, and 52. What happens when the shift is a multiple of 26? Why?

Exercise 3. Write a function `dcaesar` which will decrypt a ciphertext constructed by a Caesar's cipher. It should have the same signature as `caesar`. Try it on the examples from the previous exercise.

Exercise 4. Observe that the shift amount in a Caesar cipher is very limited: Any shift of `d` is equivalent to a shift by `d % 26`. (For instance shifting by 12 and 38 is the same thing, due to wrap around at 26.) Based on this observation, how strong do you think the Caesar's cipher is? Describe a simple attack that will recover the plaintext and automate it using Cryptol. Use your function to crack the ciphertext `JHLZHYJPWOLYPZDLHR`.

Exercise 5. One classic trick to strengthen ciphers is to use multiple keys. By repeatedly encrypting the plaintext multiple times we can hope that it will be more resistant to attacks. Do you think this scheme might make the Caesar cipher stronger?

Exercise 6. What happens if you pass `caesar` a plaintext that has non-uppercase letters in it? (Let's say a digit.) How can you fix this deficiency?

3.2 Vigenère cipher

The Vigenère cipher is a variation on the Caesar's cipher, where one uses multiple shift amounts according to a keyword [25]. Despite its simplicity, it earned the notorious description *le chiffre indéchiffrable* ("the

3.3. The atbash

indecipherable cipher” in French), as it was unbroken for a long period of time. It was very popular in the 16th century and onwards, only becoming routinely breakable by mid-19th century or so.

To illustrate the operation of the Vigenère cipher, let us consider the plaintext `ATTACKATDAWN`. The cryptographer picks a key, let’s say `CRYPTOL`. We line up the plaintext and the key, repeating the key as much as as necessary, as in the top two lines of the following:

```
Plaintext :   ATTACKATDAWN
Cipherkey :   CRYPTOLCRYPT
Ciphertext:   CKRPVYLVUYLEG
```

We then proceed pair by pair, shifting the plaintext character by the distance implied by the corresponding key character. The first pair is `A-C`. Since `C` is two positions away from `A` in the alphabet, we shift `A` by two positions, again obtaining `C`. The second pair `T-R` proceeds similarly: Since `R` is 17 positions away from `A`, we shift `T` down 17 positions, wrapping around `Z`, obtaining `K`. Proceeding in this fashion, we get the ciphertext `CKRPVYLVUYLEG`. Note how each step of the process is a simple application of the Caesar’s cipher.

Exercise 7. One component of the Vigenère cipher is the construction of the repeated key. Write a function `cycle` with the following signature:

```
cycle : {n, a} (fin n, n >= 1) => [n]a -> [inf]a
```

such that it returns the input sequence appended to itself repeatedly, turning it into an infinite sequence. Why do we need the predicate `n >= 1`?

Exercise 8. Program the Vigenère cipher in Cryptol. It should have the signature:

```
vigenere : {n, m} (fin n, n >= 1) => (String n, String m) -> String m
```

where the first argument is the key and the second is the plaintext. Note how the signature ensures that the input string and the output string will have precisely the same number of characters, `m`. (**Hint:** Use Caesar’s cipher repeatedly.)

Exercise 9. Write the decryption routine for Vigenère. Then decode `"XZETGSCGTTCMGEQGAGRDEQC"` with the key `"CRYPTOL"`.

Exercise 10. A known-plaintext attack is one where an attacker obtains a plaintext-ciphertext pair, without the key. If the attacker can figure out the key based on this pair then he can break all subsequent communication until the key is replaced. Describe how one can break the Vigenère cipher if a plaintext-ciphertext pair is known.

3.3 The atbash

The atbash cipher is a form of a shift cipher, where each letter is replaced by the letter that occupies its mirror image position in the alphabet. That is, `A` is replaced by `Z`, `B` by `Y`, etc. Needless to say the atbash is hardly worthy of cryptographic attention, as it is trivial to break.

Exercise 11. Program the atbash in Cryptol. What is the code for `ATTACKATDAWN`?

Exercise 12. Program the atbash decryption in Cryptol. Do you have to write any code at all? Break the code `ZGYZHSRHHVUUVWXIBKGRMT`.

3.4 Substitution ciphers

Substitution ciphers generalize all the ciphers we have seen so far, by allowing arbitrary substitutions to be made for individual “components” of the plaintext [24]. Note that these components need not be individual characters, but rather can be pairs or even triples of characters that appear consecutively in the text. (The multi-character approach is termed *polygraphic*.) Furthermore, there are variants utilizing multiple *polyalphabetic* mappings, as opposed to a single *monoalphabetic* mapping. We will focus on monoalphabetic simple substitutions, although the other variants are not fundamentally more difficult to implement.

Tip. For the exercises in this section we will use a running key repeatedly. To simplify your interaction with Cryptol, put the following definition in your program file:

```
substKey : String 26
substKey = "FJHWOTYRXMKBPZIAZEVNULSGDCQ"
```

The intention is that `substKey` maps A to F, B to J, C to H, and so on.

Exercise 13. Implement substitution ciphers in Cryptol. Your function should have the signature:

```
subst : {n} (String 26, String n) -> String n
```

where the first element is the key (like `substKey`). What is the code for “SUBSTITUTIONSSAVETHEDAY” for the key `substKey`?

Decryption Programming decryption is more subtle. We can no longer use the simple selection operation (`@`) on the key. Instead, we have to search for the character that maps to the given ciphertext character.

Exercise 14. Write a function `invSubst` with the following signature:

```
invSubst : (String 26, Char) -> Char
```

such that it returns the mapped plaintext character. For instance, with `substKey`, F should get you A, since the key maps A to F:

```
Cryptol> invSubst (substKey, 'F')
A
```

And similarly for other examples:

```
Cryptol> invSubst (substKey, 'J')
B
Cryptol> invSubst (substKey, 'C')
Y
Cryptol> invSubst (substKey, 'Q')
Z
```

One question is what happens if you search for a non-existing character. In this case you can just return 0, a non-valid ASCII character, which can be interpreted as *not found*.

3.5. The scytale

Hint. Use a fold (see Pg. 29).

Exercise 15. Using `invSubst`, write the decryption function `dSubst`. It should have the exact same signature as `subst`. Decrypt `FUUFHKFUWFGI`, using our running key.

Exercise 16. Try the substitution cipher with the key `AAAABBBBCCCCDDDDDEEEFFFFFGG`. Does it still work? What is special about `substKey`?

3.5 The scytale

The scytale is one of the oldest cryptographic devices ever, dating back to at least the first century A.D. [23]. Ancient Greeks used a leather strip on which they would write their plaintext message. The strip would be wrapped around a rod of a certain diameter. Once the strip is completely wound, they would read the text row-by-row, essentially transposing the letters and constructing the ciphertext. Since the ciphertext is formed by a rearrangement of the plaintext, the scytale is an example of a transposition cipher. To decrypt, the ciphertext needs to be wrapped around a rod of the same diameter, reversing the process. The cipherkey is essentially the diameter of the rod used. Needless to say, the scytale does not provide a very strong encryption mechanism.

Abstracting away from the actual rod and the leather strip, encryption is essentially writing the message column-by-column in a matrix and reading it row-by-row. Let us illustrate with the message `ATTACKATDAWN`, where we can fit 4 characters per column:

```
ACD
TKA
TAW
ATN
```

To encrypt, we read the message row-by-row, obtaining `ACDTKATAWATN`. If the message does not fit properly (i.e., if it has empty spaces in the last column), it can be padded by `z`'s or some other agreed upon character. To decrypt, we essentially reverse the process, by writing the ciphertext row-by-row and reading it column-by-column.

Notice how the scytale's operation is essentially matrix transposition. Therefore, implementing the scytale in Cryptol is merely an application of the `transpose` function. All we need to do is group the message by the correct number of elements using `split`. Below, we define the `diameter` to be the number of columns we have. The type synonym `Message` ensures we only deal with strings that properly fit the "rod," by using `r` number of rows:

```
scytale : {row, diameter} (fin row, fin diameter)
  => String (row * diameter) -> String (diameter * row)
scytale msg = join (transpose msg')
  where   msg' : [diameter][row][8]
         msg' = split msg
```

The signature on `msg'` is revealing: We are taking a string that has `diameter * row` characters in it, and chopping it up so that it has `row` elements, each of which is a string that has `diameter` characters in it. Here is Cryptol in action, encrypting the message `ATTACKATDAWN`:

```
Cryptol> :set ascii=on
Cryptol> scytale "ATTACKATDAWN"
"ACDTKATAWATN"
```

Decryption is essentially the same process, except we have to `split` so that we get `diameter` elements out:

```
dScytale : {row, diameter} (fin row, fin diameter)
          => String (row * diameter) -> String (diameter * row)
dScytale msg = join (transpose msg')
  where msg' : [row][diameter][8]
        msg' = split msg
```

Again, the type on `msg'` tells Cryptol that we now want `diameter` strings, each of which is `row` long. It is important to notice that the definitions of `scytale` and `dscytale` are precisely the same, except for the signature on `msg'`! When viewed as a matrix, the types precisely tell which transposition we want at each step. We have:

```
Cryptol> dScytale "ACDTKATAWATN"
"ATTACKATDAWN"
```

Exercise 17. What happens if you comment out the signature for `msg'` in the definition of `scytale`? Why?

Exercise 18. How would you attack a scytale encryption, if you don't know what the diameter is?

Chapter 4

The Enigma machine

The Enigma machine is probably the most famous of all cryptographic devices in history, due to the prominent role it played in WWII [18]. The first Enigma machines were available around 1920s, with various models in the market for commercial use. When Germans used the Enigma during WWII, they were using a particular model referred to as the *Wehrmacht Enigma*, a fairly advanced model available at the time.

The most important role of Enigma is in its role in the use of automated machines to aid in secret communication, or what is known as *mechanizing secrecy*. One has to understand that computers as we understand them today were not available when Enigma was in operation. Thus, the Enigma employed a combination of mechanical (keyboard, rotors, etc.) and electrical parts (lamps, wirings, etc.) to implement its functionality. However, our focus in this chapter will not be on the mechanical aspects of Enigma at all. For a highly readable account of that, we refer the reader to Singh's excellent book on cryptography [15]. Instead, we will model Enigma in Cryptol in an algorithmic sense, implementing Enigma's operations without any reference to the underlying mechanics. More precisely, we will model an Enigma machine that has a plugboard, three interchangeable scramblers, and a fixed reflector.

4.1 The plugboard

Enigma essentially implements a polyalphabetic substitution cipher (Section 3.4), consisting of a number of rotating units that jumble up the alphabet. The first component is the so called plugboard (*steckerbrett* in German). In the original Enigma, the plugboard provided a means of interchanging 6-pairs of letters. For instance, the plugboard could be set-up so that pressing the **B** key would actually engage the **Q** key, etc. We will slightly generalize and allow any number of pairings, as we are not limited by the availability of cables or actual space to put them in a box! Viewed in this sense, the plugboard is merely a permutation of the alphabet. In Cryptol, we can represent the plugboard combination by a string of 26 characters, corresponding to the pairings for each letter in the alphabet from A to Z:

```
type Permutation = String 26
type Plugboard = Permutation
```

For instance, the plugboard matching the pairs A-H, C-G, Q-X, T-V, U-Y, W-M, and O-L can be created as follows:

```
plugboard : Plugboard
plugboard = "HBGDEFCAIJKOWNLPXRSVYTMQUZ"
```

Note that if a letter is not paired through the plugboard, then it goes untouched, i.e., it is paired with itself.

Exercise 1. Use Cryptol to verify that the above plugboard definition indeed implements the pairings we wanted.

Note. In Enigma, the plugboard pairings are symmetric; if A maps to H, then H must map to A.

4.2 Scrambler rotors

The next component of the Enigma are the rotors that scramble the letters. Rotors (*walzen* in German) are essentially permutations, with one little twist: as their name implies, they rotate. This rotation ensures that the next character the rotor will process will be encrypted using a different alphabet, thus giving Enigma its polyalphabetic nature.

The other trick employed by Enigma is how the rotations are done. In a typical setup, the rotors are arranged so that the first rotor rotates at every character, while the second rotates at every 26th, the third at every 676th ($= 26 * 26$), etc. In a sense, the rotors work like the odometer in your car, one full rotation of the first rotor triggers the second, whose one full rotation triggers the third, and so on. In fact, more advanced models of Enigma allowed for two notches per rotor, i.e., two distinct positions on the rotor that will allow the next rotor in sequence to rotate itself. We will allow ourselves to have any number of notches, by simply pairing each substituted letter with a bit value saying whether it has an associated notch:¹

```
type Rotor = [26](Char, Bit)
```

The function `mkRotor` will create a rotor for us from a given permutation of the letters and the notch locations:²

```
mkRotor : {n} (fin n) => (Permutation, String n) -> Rotor
mkRotor (perm, notchLocations) = [ (p, elem (p, notchLocations))
                                   | p <- perm
                                   ]
```

Let us create a couple of rotors with notches:

```
rotor1, rotor2, rotor3 : Rotor
rotor1 = mkRotor ("RJICAWVQZODLUPYFEHXSMTKNGB", "IO")
rotor2 = mkRotor ("DWYOLETKNVQPHURZJMSFIGXCBA", "B")
rotor3 = mkRotor ("FGKMAJWUOVNRYIZETDPSHBLCQX", "CK")
```

For instance, `rotor1` maps A to R, B to J, . . . , and Z to B in its initial position. It will engage its notch if one of the permuted letters I or O appear in its first position.

Exercise 2. Write out the encrypted letters for the sequence of 5 C's for `rotor1`, assuming it rotates in each step. At what points does it engage its own notch to signal the next rotor to rotate?

¹The type definition for `Char` was given in Example 3.4-14.

²The function `elem` was defined in Exercise 2.13-59.

4.3 Connecting the rotors: notches in action

The original Enigma had three interchangeable rotors. The operator chose the order they were placed in the machine. In our model, we will allow for an arbitrary number of rotors. The tricky part of connecting the rotors is ensuring that the rotations of each are done properly.

Let us start with a simpler problem. If we are given a rotor and a particular letter to encrypt, how can we compute the output letter and the new rotor position? First of all, we will need to know if the rotor should rotate itself, that is if the notch between this rotor and the previous one was activated. Also, we need to find out if the act of rotation in this rotor is going to cause the next rotor to rotate. We will model this action using the Cryptol function `scramble`:

```
scramble : (Bit, Char, Rotor) -> (Bit, Char, Rotor)
```

The function `scramble` takes a triple `(rotate, c, rotor)`:

- `rotate`: if `True`, this rotor will rotate before encryption. Indicates that the notch between this rotor and the previous one was engaged,
- `c`: the character to encrypt, and
- `rotor`: the current state of the rotor.

Similarly, the output will also be a triple:

- `notch`: `True` if the notch on this rotor engages, i.e., if the next rotor should rotate itself,
- `c'`: the result of encrypting (substituting) for `c` with the current state of the rotor.
- `rotor'`: the new state of the rotor. If no rotation was done this will be the same as `rotor`. Otherwise it will be the new substitution map obtained by rotating the old one to the left by one.

Coding `scramble` is straightforward:

```
scramble (rotate, c, rotor) = (notch, c', rotor')
  where
    (c', _)    = rotor @ (c - 'A')
    (_, notch) = rotor @ 0
    rotor'    = if rotate then rotor <<< 1 else rotor
```

To determine `c'`, we use the substitution map to find out what this rotor maps the given character to, with respect to its current state. Note how Cryptol's pattern matching notation helps with extraction of `c'`, as we only need the character, not whether there is a notch at that location. (The underscore character use, `'_'`, means that we do not need the value at the position, and hence we do not give it an explicit name.) To determine if we have our notch engaged, all we need to do is to look at the first element's notch value, using Cryptol's selection operator (`@ 0`), and we ignore the permutation value there this time, again using pattern matching. Finally, to determine `rotor'` we merely rotate-left by 1 if the `rotate` signal was received. Otherwise, we leave the `rotor` unchanged.

Exercise 3. Redo Exercise 2, this time using Cryptol and the `scramble` function.

Note. The actual mechanics of the Enigma machine were slightly more complicated: due to the keyboard mechanism and the way notches were mechanically built, the first rotor was actually rotating before the encryption took place. Also, the middle rotor could double-step if it engages its notch right after the third rotor does [2].

We will take the simpler view here and assume that each key press causes an encryption to take place, *after* which the rotors do their rotation, getting ready for the next input. The mechanical details, while historically important, are not essential for our modeling purposes here. Also, the original Enigma had *rings*, a relatively insignificant part of the whole machine, that we ignore here.

Sequencing the rotors Now that we have the rotors modeled, the next task is to figure out how to connect them in a sequence. As we mentioned, Enigma had 3 rotors originally (later versions allowing 4). The three rotors each had a single notch (later versions allowing double notches). Our model allows for arbitrary number of rotors and arbitrary number of notches on each. The question we now tackle is the following: Given a sequence of rotors, how do we run them one after the other? We are looking for a function with the following signature:

```
joinRotors : {n} (fin n) => ([n]Rotor, Char) -> ([n]Rotor, Char)
```

That is, we receive n rotors and the character to be encrypted, and return the updated rotors (accounting for their rotations) and the final character. The implementation is an instance of the fold pattern (Section 2.13), using the `scramble` function we have just defined:

```
joinRotors (rotors, inputChar) = (rotors', outputChar)
  where
    initRotor = mkRotor (['A' .. 'Z'], [])
    ncrs : [n+1](Bit, [8], Rotor)
    ncrs = [(True, inputChar, initRotor)
            # [ scramble (notch, char, r)
              | r <- rotors
              | (notch, char, rotor') <- ncrs
            ]
    rotors' = tail [ r | (_, _, r) <- ncrs ]
    (_, outputChar, _) = ncrs ! 0
```

The workhorse in `joinRotors` is the definition of `ncrs`, a mnemonic for *notches-chars-rotors*. The idea is fairly simple. We simply iterate over all the given rotors (`r <- rotors`), and `scramble` the current character `char`, using the rotor `r` and the notch value `notch`. These values come from `ncrs` itself, using the fold pattern encoded by the comprehension. The only question is what is the seed value for this fold?

The seed used in `ncrs` is `(True, inputChar, initRotor)`. The first component is `True`, indicating that the very first rotor should always rotate itself at every step. The second element is `inputChar`, which is the input to the whole sequence of rotors. The only mysterious element is the last one, which we have specified as `initRotor`. This rotor is defined so that it simply maps the letters to themselves with no notches on it, by a call to the `mkRotor` function we have previously defined. This rotor is merely a place holder to kick off the computation of `ncrs`, it acts as the identity element in a sequence of rotors. To compute `rotors'`, we merely project the third component of `ncrs`, being careful about skipping the first element using `tail`. Finally, `outputChar` is merely the output coming out of the final rotor, extracted using `!0`. Note how we use Cryptol's pattern matching to get the second component out of the triple in the last line.

4.4. The reflector

Exercise 4. Is the action of `initRotor` ever used in the definition of `joinRotors`?

Exercise 5. What is the final character returned by the expression:

```
joinRotors ([rotor1 rotor2 rotor3], 'F')
```

Use paper and pencil to figure out the answer by tracing the execution of `joinRotors` before running it in Cryptol!

4.4 The reflector

The final piece of the Enigma machine is the reflector (*umkehrwalze* in German). The reflector is another substitution map. Unlike the rotors, however, the reflector did not rotate. Its main function was ensuring that the process of encryption was reversible: The reflector did one final jumbling of the letters and then sent the signal back through the rotors in the *reverse* order, thus completing the loop and allowing the signal to reach back to the lamps that would light up. For our purposes, it suffices to model it just like any other permutation:

```
type Reflector = Permutation
```

Here is one example:

```
reflector : Reflector
reflector = "FEIPBATSCYVUWZQDOXHGLKMRJN"
```

Like the plugboard, the reflector is symmetric: If it maps `B` to `E`, it should map `E` to `B`, as in the above example. Furthermore, the Enigma reflectors were designed so that they never mapped any character to themselves, which is true for the above permutation as well. Interestingly, this idea of a non-identity reflector (i.e., never mapping any character to itself) turned out to be a weakness in the design, which the allies exploited in breaking the Enigma during WWII [15].

Exercise 6. Write a function `checkReflector` with the signature:

```
checkReflector : Reflector -> Bit
```

such that it returns `True` if a given reflector is good (i.e., symmetric and non-self mapping) and `False` otherwise. Check that our definition of `reflector` above is a good one. (**Hint:** Use the `all` function you have defined in Exercise 2.9-52.)

4.5 Putting the pieces together

We now have all the components of the Enigma: the plugboard, rotors, and the reflector. The final task is to implement the full loop. The Enigma ran all the rotors in sequence, then passed the signal through the reflector, and ran the rotors in reverse one more time before delivering the signal to the lamps.

Before proceeding, we will define the following two helper functions:

```
substFwd, substBwd : (Permutation, Char) -> Char
substFwd (perm, c) = perm @ (c - 'A')
substBwd (perm, c) = invSubst (perm, c)
```

(You have defined the `invSubst` function in Exercise 3.4-14.) The `substFwd` function simply returns the character that the given permutation, whether from the plugboard, a rotor, or the reflector. Conversely, `substBwd` returns the character that the given permutation maps *from*, i.e., the character that will be mapped to `c` using the permutation.

Exercise 7. Using Cryptol, verify that `substFwd` and `substBwd` return the same elements for each letter in the alphabet for `rotor1`.

Exercise 8. Show that `substFwd` and `substBwd` are exactly the same operations for the reflector. Why?

The route back One crucial part of the Enigma is the running of the rotors in reverse after the reflector. Note that this operation ignores the notches, i.e., the rotors do not turn while the signal is passing the second time through the rotors. (In a sense, the rotations happen after the signal completes its full loop, getting to the reflector and back.) Consequently, it is much easier to code as well (compare this code to `joinRotors`, defined in Section 4.3):

```
backSignal : {n} (fin n) => ([n]Rotor, Char) -> Char
backSignal (rotors, inputChar) = cs ! 0
  where
    cs = [inputChar] # [ substBwd ([ p | (p, _) <- r ], c)
                        | r <- reverse rotors
                        | c <- cs
                        ]
```

Note that we explicitly reverse the rotors in the definition of `cs`. (The definition of `cs` is another typical example of a fold. See Pg. 29.)

Given all this machinery, coding the entire Enigma loop is fairly straightforward:

```
//enigmaLoop : {n} (fin n) => (Plugboard, [n]Rotor, Reflector, Char)
//                                     -> ([n]Rotor, Char)
enigmaLoop (pboard, rotors, refl, c0) = (rotors', c5)
  where
    // 1. First run through the plugboard
    c1 = substFwd (pboard, c0)
    // 2. Now run all the rotors forward
    (rotors', c2) = joinRotors (rotors, c1)
    // 3. Pass through the reflector
    c3 = substFwd (refl, c2)
    // 4. Run the rotors backward
    c4 = backSignal(rotors, c3)
    // 5. Finally, back through the plugboard
    c5 = substBwd (pboard, c4)
```

4.6 The state of the machine

We are almost ready to construct our own Enigma machine in Cryptol. Before doing so, we will take a moment to represent the state of the Enigma machine as a Cryptol record, which will simplify our final

4.7. Encryption and decryption

construction. At any stage, the state of an Enigma machine is given by the status of its rotors. We will use the following record to represent this state, for an Enigma machine containing n rotors:

```
type Enigma n = { plugboard : Plugboard,
                 rotors    : [n]Rotor,
                 reflector : Reflector
               }
```

To initialize an Enigma machine, the operator provides the plugboard settings, rotors, the reflector. Furthermore, the operator also gives the initial positions for the rotors. Rotors can be initially rotated to any position before put together into the machine. We can capture this operation with the function `mkEnigma`:

```
mkEnigma : {n} (Plugboard, [n]Rotor, Reflector, [n]Char)
          -> Enigma n
mkEnigma (pboard, rs, refl, startingPositions) =
  { plugboard = pboard,
    rotors    = [ r <<< (s - 'A')
                | r <- rs
                | s <- startingPositions
                ],
    reflector = refl
  }
```

Note how we rotate each given rotor to the left by the amount given by its starting position.

Given this definition, let us construct an Enigma machine out of the components we have created so far, using the starting positions GCR for the rotors respectively:

```
modelEnigma : Enigma 3
modelEnigma = mkEnigma (plugboard, [rotor1, rotor2, rotor3],
                      reflector, "GCR")
```

We now have an operational Enigma machine coded up in Cryptol!

4.7 Encryption and decryption

Equipped with all the machinery we now have, coding Enigma encryption is fairly straightforward:

```
enigma : {n, m} (fin n, fin m) => (Enigma n, String m) -> String m
enigma (m, pt) = tail [ c | (_, c) <- rcs ]
  where rcs = [(m.rotors, '*')] #
             [ enigmaLoop (m.plugboard, r, m.reflector, c)
             | c <- pt
             | (r, _) <- rcs
             ]
```

The function `enigma` takes a machine with n rotors and a plaintext of m characters, returning a ciphertext of m characters back. It is yet another application of the fold pattern, where we start with the initial set of rotors and the placeholder character `*` (which could be anything) to seed the fold. Note how the change in rotors

is reflected in each iteration of the fold, through the `enigmaLoop` function. At the end, we simply drop the rotors from `rcs`, and take the `tail` to skip over the seed character `*`.

Here is our Enigma in operation:

```
Cryptol> :set ascii=on
Cryptol> enigma (modelEnigma, "ENIGMAWASAREALLYCOOLMACHINE")
"UPEKTBSDROBVTUJGNCEHHGBXGTF"
```

Decryption As we mentioned before, Enigma was a self-decrypting machine, that is, encryption and decryption are precisely the same operations. Thus, we can define:

```
dEnigma : {n, m} (fin n, fin m) => (Enigma n, String m) -> String m
dEnigma = enigma
```

And decrypt our above message back:

```
Cryptol> dEnigma (modelEnigma, "UPEKTBSDROBVTUJGNCEHHGBXGTF")
"ENIGMAWASAREALLYCOOLMACHINE"
```

We have successfully performed our first Enigma encryption!

Exercise 9. Different models of Enigma came with different sets of rotors. You can find various rotor configurations on the web [19]. Create models of these rotors in Cryptol, and run sample encryptions through them.

Exercise 10. As we have mentioned before, Enigma implements a polyalphabetic substitution cipher, where the same letter gets mapped to different letters during encryption. The period of a cipher is the number of characters before the encryption repeats itself, mapping the same sequence of letters in the plaintext to the same sequence of letters in the ciphertext. What is the period of an Enigma machine with n rotors?

Exercise 11. Construct a string of the form `CRYPTOLXXX...XCRYPTOL`, where `...`'s are filled with enough number of `x`'s such that encrypting it with our `modelEnigma` machine will map the instances of `"CRYPTOL"` to the same ciphertext. How many `x`'s do you need? What is the corresponding ciphertext for `"CRYPTOL"` in this encryption?

The code You can see all the Cryptol code for our Enigma simulator in Appendix C.

Chapter 5

High-assurance programming

Writing correct software is the holy grail of programming. Bugs inevitably exist, however, even in thoroughly tested projects. One fundamental issue is the lack of support in typical programming languages to let the programmer *state* what it means to be correct, let alone formally establish any notion of correctness. To address this shortcoming, Cryptol advocates the high-assurance programming approach: programmers explicitly state correctness properties along with their code, which are explicitly checked by the Cryptol toolset. Properties are not comments or mere annotations, so there is no concern that they will become obsolete as your code evolves. The goal of this chapter is to introduce you to these tools, and to the notion of high-assurance programming in Cryptol via examples.

5.1 Writing properties

Consider the equality:

$$x^2 - y^2 = (x - y) * (x + y)$$

Let us write two Cryptol functions that capture both sides of this equation:

```
sqDiff1 (x, y) = x^2 - y^2
sqDiff2 (x, y) = (x-y) * (x+y)
```

We would like to express the property that `sqDiff1` and `sqDiff2` are precisely the same functions: Given the same `x` and `y`, they should return exactly the same answer. We can express this property in Cryptol using a `properties` declaration:

```
sqDiffsCorrect : ([8], [8]) -> Bit
property sqDiffsCorrect (x, y) = sqDiff1 (x, y) == sqDiff2 (x, y)
```

The above declaration reads as follows: `sqDiffsCorrect` is a property stating that for all `x` and `y`, the expression `sqDiff1 (x, y) == sqDiff2(x, y)` evaluates to `True`. Furthermore, the type signature restricts the type of the property to apply to only 8-bit values. As usual, the type-signature is optional. If not given, Cryptol will infer one for you.

Exercise 1. Put the above property in a file and load it into Cryptol. Then issue the command:

```
Cryptol> :t sqDiffsCorrect
```

What do you see?

Note. It is important to emphasize that the mathematical equality above and the Cryptol property are *not* stating precisely the same property. Remember that all Cryptol arithmetic is modular, while the mathematical equation is over arbitrary numbers, including negative, real, or even complex numbers. The takeaway of this discussion is that we are only using this example for illustration purposes: Cryptol properties relate to Cryptol programs, and should not be used for expressing mathematical theorems (unless, of course, you are stating group theory theorems or theorems in an appropriate algebra)! In particular, `sqDiffsCorrect` is a property about the Cryptol functions `sqDiff1` and `sqDiff2`, not about the mathematical equation that inspired it.

Exercise 2. Write a property `revRev` stating that reverse of a reverse returns a sequence unchanged.

Exercise 3. Write a property `appAssoc` stating that append is an associative operator.

Exercise 4. Write a property `revApp` stating that appending two sequences and reversing the result is the same as reversing the sequences and appending them in the reverse order, as illustrated in the following expression:

```
reverse ("HELLO" # "WORLD") == reverse "WORLD" # reverse "HELLO"
```

Exercise 5. Write a property `1shMul` stating that shifting left by k is the same as multiplying by 2^k .

NB. A property declaration simply introduces a property about your program, which may or may not actually hold. It is an assertion about your program, without any claim of correctness. In particular, you can clearly write properties that simply do not hold:

```
property bogus x = x != x
```

It is important to distinguish between *stating* a property and actually *proving* it. So far, our focus is purely on specification. We will focus on actual proofs in Section 5.2.

5.1.1 Property-function correspondence

In Cryptol, properties can be used just like ordinary definitions:

```
Cryptol> sqDiffsCorrect (3, 5)
True
Cryptol> :t sqDiffsCorrect
sqDiffsCorrect : ([8],[8]) -> Bit
```

That is, a property over (x, y) is the same as a function over the tuple (x, y) . We call this the property-function correspondence. Property declarations, aside from the slightly different syntax, are *precisely* the same as Cryptol functions whose return type is `Bit`. There is no separate language for writing or working with properties. We simply use the full Cryptol language write both the programs and the properties that they satisfy.

5.1. Writing properties

5.1.2 Capturing test vectors

One nice application of Cryptol properties is in capturing test vectors:

```
property inctest = [ f x == y | (x, y) <- testVector ] == ~zero
  where f x = x + 1
        testVector = [(3, 4), (4, 5), (12, 13), (78, 79)]
```

Notice that the property `inctest` does not have any parameters (no *forall* section), and thus acts as a simple `Bit` value that will be true precisely when the given test case succeeds.

5.1.3 Polymorphic properties

Just like functions, Cryptol properties can be polymorphic as well. If you want to write a property for a polymorphic function, for instance, your properties will naturally be polymorphic too. Here is a simple trivial example:

```
property multShift x = x * 2 == x << 1
```

If we ask Cryptol the type of `multShift`, we get:

```
Cryptol> :t multShift
multShift : {a} (fin a, a >= 2) => [a] -> Bit
```

That is, it is a property about all words of size at least two. The question is whether this property does indeed hold? In the particular case of `multShift` that is indeed the case, below are some examples using the property-function correspondence:

```
Cryptol> multShift (5 : [8])
True
Cryptol> multShift (5 : [10])
True
Cryptol> multShift (5 : [16])
True
```

However, this is *not* always the case for all polymorphic Cryptol properties! The following example demonstrates:

```
property flipNeverIdentity x = x != ~x
```

The property `flipNeverIdentity` states that complementing the bits of a value will always result in a different value: a property we might expect to hold intuitively. Here is the type of `flipNeverIdentity`:

```
Cryptol> :t flipNeverIdentity
flip : {a} (fin a) => a -> Bit
```

So, the only requirement on `flipNeverIdentity` is that it receives some finite type. Let us try some examples:

```
Cryptol> flipNeverIdentity True
True
Cryptol> flipNeverIdentity 3
True
Cryptol> flipNeverIdentity [1 2]
True
```

However:

```
Cryptol> flipNeverIdentity (0 : [0])
False
```

That is, when given a 0-bit wide value, the complement will in fact do nothing and return its argument unchanged! Therefore, the property `flipNeverIdentity` is not valid, since it holds at certain monomorphic types, but not at all types.

Exercise 6. Demonstrate another monomorphic type where `flipNeverIdentity` does *not* hold.

NB. The moral of this discussion is that the notion of polymorphic validity (i.e., that a given polymorphic property will either hold at all of its monomorphic instances or none) does not hold in Cryptol. A polymorphic property can be valid at some, all, or no instances of it.

Exercise 7. The previous exercise might lead you to think that it is the 0-bit word type (`[0]`) that is at the root of the polymorphic validity issue. This is not true. Consider the following example:

```
property widthPoly x = (w == 15) || (w == 531)
  where w = width x
```

What is the type of `widthPoly`? At what instances does it hold? Write a property `evenWidth` that holds only at even-width word instances.

5.2 Establishing correctness

Our focus so far has been using Cryptol to *state* properties of our programs, without actually trying to prove them correct. This separation of concerns is essential for a pragmatic development approach. Properties act as contracts that programmers state along with their code, which can be separately checked by the toolset [7]. This approach allows you to state the properties you want, and then work on your code until the properties are indeed satisfied. Furthermore, properties stay with your program forever, so they can be checked at a later time to ensure changes (improvements/additions/optimizations etc.) did not violate the stated properties.

5.2.1 Formal proofs

Recall our very first property, `sqDiffsCorrect`, from Section 5.1. We will now use Cryptol to actually prove it automatically. To prove `sqDiffsCorrect`, use the command `:prove`:

```
Cryptol> :prove sqDiffsCorrect
Q.E.D.
```

5.2. Establishing correctness

Note that the above might take a while to complete, as a formal proof is being produced behind the scenes. Once Cryptol formally establishes the property holds, it prints “Q.E.D.” to tell the user the proof is complete.

NB. Cryptol uses off-the-shelf SAT and SMT solvers to perform these formal proofs [7]. By default, Cryptol will use the CVC4 SMT solver [16] under the hood, but it can be configured to use other SAT/SMT solvers as well, such as SRI’s Yices [26], or Microsoft Research’s Z3 [17]¹. Note that the `:prove` command is a push-button tool: once the proof starts there is no user involvement. Of course, the external tool used may not be able to complete all the proofs in a feasible amount of time, naturally.

5.2.2 Counterexamples

Of course, properties can very well be invalid, due to bugs in code or the specifications themselves. In these cases, Cryptol will always print a counterexample value demonstrating why the property does not hold. Here is an example demonstrating what happens when things go wrong:

```
failure : [8] -> Bit
property failure x = x == x+1
```

We have:

```
Cryptol> :prove failure
failure 0 = False
```

Cryptol tells us that the property is falsifiable, and then demonstrates a particular value (0 in this case) that it fails at. These counterexamples are extremely valuable for debugging purposes.

If you try to prove an invalid property that encodes a test vector (Section 5.1.2), then you will get a mere indication that you have a contradiction, since there is no universally quantified variables to instantiate to show you a counterexample. If the expression evaluates to `True`, then it will be a trivial proof, as expected:

```
Cryptol> :prove False
False = False
Cryptol> :prove True
Q.E.D.
Cryptol> :prove 2 == 3
2==3 = False
Cryptol> :prove reverse [1, 2] == [1, 2]
reverse [1, 2] == [1,2] = False
Cryptol> :prove 1+1 == 0
Q.E.D.
```

The very last example demonstrates modular arithmetic in operation, as usual.

5.2.3 Dealing with polymorphism

As we mentioned before, Cryptol properties can be polymorphic. As we explored before, we cannot directly prove polymorphic properties as they may hold for certain monomorphic instances while not for others. In

¹To do this, first install the package(s) from the URLs provided in the bibliography. Once a prover has been installed you can activate it with, for example, `:set prover=z3`.

this cases, we must tell Cryptol what particular monomorphic instance we would like it to prove the property at. Let us demonstrate this with the `multShift` property from Section 5.1.3:

```
Cryptol> :prove multShift
Not a monomorphic type:
{a} (a >= 2, fin a) => [a] -> Bit
```

Cryptol is telling us that it cannot prove a polymorphic property directly. We can, however, give a type annotation to monomorphise it, and then prove it at a desired instance:

```
Cryptol> :prove multShift : [16] -> Bit
Q.E.D.
```

In fact, you can use this very same technique to pass any bit-valued function to the `:prove` command:

```
Cryptol> :prove dbl where dbl x = (x:[8]) * 2 == x+x
Q.E.D.
```

Of course, a λ -expression (Section 2.12.2) would work just as well too:

```
Cryptol> :prove \x -> (x:[8]) * 2 == x+x
Q.E.D.
```

Exercise 8. Prove the property `revRev` you wrote in Exercise 5.1-2. Try different monomorphic instantiations.

Exercise 9. Prove the property `appAssoc` you wrote in Exercise 5.1-3, at several different monomorphic instances.

Exercise 10. Prove the property `revApp` you wrote in Exercise 5.1-4, at several different monomorphic instances.

Exercise 11. Prove the property `lshMul` you wrote in Exercise 5.1-5, at several different monomorphic instances.

Exercise 12. Use the `:prove` command to prove and demonstrate counterexamples for the property `widthPoly` defined in Exercise 5.1-7, using appropriate monomorphic instances.

5.2.4 Conditional proofs

It is often the case that we are interested in a property that only holds under certain conditions. For instance, in Exercise 2.10-36 we have explored the relationship between Cryptol's division, multiplication, and modulus operators, where we asserted the following property:

$$x = (x/y) \times y + (x \% y)$$

Obviously, this relationship holds only when $y \neq 0$. The idea behind a conditional Cryptol property is that we would like to capture these side-conditions formally in our specifications.

We simply use an ordinary `if-then-else` expression in Cryptol to write conditional properties (at least until we add boolean logic operators to Cryptol). If the condition is invalid, we simply return `True`, indicating that we are not interested in that particular case. Depending on how natural it is to express the side-condition or its negation, you can use one of the following two patterns:

5.3. Automated random testing

```
if side-condition-holds      if side-condition-fails
then property-expression    then True
else True                    else property-expression
```

Exercise 13. Express the relationship between division, multiplication, and modulus using a conditional Cryptol property. Prove the property for various monomorphic instances.

Recognizing messages Our work on classic ciphers (Chapter 3) and the enigma (Chapter 4) involved working with messages that contained the letters 'A' .. 'Z' only. When writing properties about these ciphers it will be handy to have a recognizer for such messages, as we explore in the next exercise.

Exercise 14. Write a function:

```
validMessage : {n} (fin n) => String n -> Bit
```

that returns `True` exactly when the input only consists of the letters 'A' through 'Z'. (**Hint:** Use the functions `all` defined in Exercise 2.9-52, and `elem` defined in Exercise 2.13-59.)

Exercise 15. Recall the pair of functions `caesar` and `dCaesar` from Section 3.1. Write a property, named `caesarCorrect`, stating that `caesar` and `dCaesar` are inverses of each other for all `d` (shift amount) and `msg` (message)'s. Is your property valid? What extra condition do you need to assert on `msg` for your property to hold? Prove the property for all messages of length 10.

Exercise 16. Write and prove a property for the `modelEnigma` machine (Page 49), relating the `enigma` and `dEnigma` functions from Section 4.7.

This may take a long time to prove, depending on the speed of your machine, and the prover you choose.

5.3 Automated random testing

Cryptol's `:prove` command constructs rigorous formal proofs using push-button tools.² The underlying technique used by Cryptol (SAT- and SMT-based equivalence checking) is complete, i.e., it will always either prove the property or find a counterexample. In the worst case, however, the proof process might take infeasible amounts of resources, potentially running out of memory or taking longer than the amount of time you are willing to wait.

What is needed for daily development tasks is a mechanism to gain some confidence on the correctness of the properties without paying the price of formally proving them. This is the goal of Cryptol's `:check` command, inspired by Haskell's quick-check library [3]. Instead of trying to formally prove your property, `:check` tests it at random values to give you quick feedback. This approach is very suitable for rapid development. By using automated testing you get frequent and quick updates from Cryptol regarding the status of your properties, as you work through your code. If you introduce a bug, it is likely (although not guaranteed) that the `:check` command will alert you right away. Once you are satisfied with your code, you can use the `:prove` command to conduct the formal proofs, potentially leaving them running overnight.

The syntax of the `:check` command is precisely the same as the `:prove` command. By default, it will run your property over 100 randomly generated test cases.

²While some of the solvers that Cryptol uses are capable of *emitting* proofs, such functionality is not exposed as a Cryptol feature.

Exercise 17. Use the `:check` command to test the property `caesarCorrect` you have defined in Exercise 5.2.4-14, for messages of length 10. Use the command `:set tests=1000` to change the number of test cases to 1,000. Observe the test coverage statistics reported by Cryptol. How is the total number of cases computed?

Exercise 18. If the property is *small* in size, `:check` might as well prove/disprove it. Try the following commands:

```
:check True
:check False
:check \x -> x==(x:[8])
```

Exercise 19. Write a bogus property that will be very easy to disprove using `:prove`, while `:check` will have a hard time obtaining the counterexample. The moral of this exercise is that you should try `:prove` early in your development and not get too comfortable with the results of `:check`!

Bulk operations If you use `:check` and `:prove` commands without any arguments, Cryptol will check and prove all the properties defined in your program. This is a simple means of exercising all your properties automatically.

5.4 Checking satisfiability

Closely related to proving properties is the notion of checking satisfiability. In satisfiability checking, we would like to find arguments to a bit-valued function such that it will evaluate to `True`, i.e., it will be satisfied.

One way to think about satisfiability checking is *intelligently* searching for a solution. Here is a simple example. Let us assume we would like to compute the modular square-root of 9 as an 8-bit value. The obvious solution is 3, of course, but we are wondering if there are other solutions to the equation $x^2 \equiv 9 \pmod{2^8}$. To get started, let us first define a function that will return `True` if its argument is a square-root of 9:

```
isSqrtOf9 : [8] -> Bit
isSqrtOf9 x = x*x == 9
```

Any square-root of 9 will make the function `isSqrtOf9` return `True`, i.e., it will *satisfy* it. Thus, we can use Cryptol's satisfiability checker to find those values of `x` automatically:

```
Cryptol> :sat isSqrtOf9
isSqrtOf9 3 = True
```

Not surprisingly, Cryptol told us that 3 is one such value. We can search for other solutions by explicitly disallowing 3:

```
Cryptol> :sat \x -> isSqrtOf9 x && ~(elem (x, [3]))
\x -> isSqrtOf9 x && ~(elem (x, [3])) 125 = True
```

Note the use of the λ -expression to indicate the new constraint. (Of course, we could have defined another function `isSqrtOf9ButNot3` for the same effect, but the λ -expression is really handy in this case.) We have used the function `elem` you have defined in Exercise 2.13-59 to express the constraint `x` must not be 3. In response, Cryptol told us that 125 is another solution. Indeed $125 * 125 = 9 \pmod{2^7}$, as you can verify separately. We can search for more:

5.4. Checking satisfiability

```
Cryptol> :sat \x -> isSqrtOf9 x && ~(elem (x, [3 125]))
\x -> isSqrtOf9 x && ~(elem (x, [3, 125]) ) 131 = True
```

And more:

```
Cryptol> :sat \x -> isSqrtOf9 x && ~(elem (x, [3 131 125]))
\x -> isSqrtOf9 x && ~(elem (x, [3, 131, 125]) ) 253 = True
```

If we try one more time, we get:

```
Cryptol> :sat \x -> isSqrtOf9 x && ~(elem (x, [3, 131, 125, 253]))
Unsatisfiable.
```

So, we can rest assured that there are exactly four 8-bit square roots of 9; namely 3, 131, 125, and 253. (Note that Cryptol can return the satisfying solutions in any order depending on the backend-solver and other configurations. What is guaranteed is that you will get precisely the same set of solutions at the end.)

The whole point of the satisfiability checker is to be able to quickly search for particular values that are solutions to potentially complicated bit-valued functions. In this sense, satisfiability checking can also be considered as an automated way to invert a certain class of functions, going back from results to arguments. Of course, this search is not done blindly, but rather using SAT and SMT solvers to quickly find the satisfying values. Cryptol's `:sat` command hides the complexity, allowing the user to focus on the specification of the problem.

Exercise 20. Fermat's last theorem states that there are no integer solutions to the equation $a^n + b^n = c^n$ when $a, b, c > 0$, and $n > 2$. We cannot code Fermat's theorem in Cryptol since we do not have arbitrary integers, but we can code the modular version of it where the exponentiation and addition is done modulo a fixed bit-size. Write a function `modFermat` with the following signature:

```
type Quad a = ([a], [a], [a], [a])
modFermat : {s} (fin s, s >= 2) => Quad s -> Bit
```

such that `modFermat (a, b, c, n)` will return `True` if the modular version of Fermat's equation is satisfied by the values of `a`, `b`, `c`, and `n`. Can you explain why you need the constraints `fin s` and `s >= 2`?

Exercise 21. Use the `:sat` command to see if there are any satisfying values for the modular version of Fermat's last theorem for various bit sizes. Surprised? What can you conclude from your observations?

5.4. Checking satisfiability

Chapter 6

AES: The Advanced Encryption Standard

AES is a symmetric key encryption algorithm (a symmetric cipher, per the discussion in chapter 3), based on the Rijndael algorithm designed by Joan Daemen and Vincent Rijmen [5]. (The term *symmetric key* means that the algorithm uses the same key for encryption and decryption.) AES was adopted in 2001 by the US government, deemed suitable for protecting classified information up to *secret* level for the key size 128, and up to the *top-secret* level for key sizes 192 and 256.

In this chapter, we will program AES in Cryptol. Our emphasis will be on clarity, as opposed to efficiency, and we shall follow the NIST standard description of AES fairly closely [13]. Referring to the standard as you work your way through this chapter is recommended.

Some surprises may be at hand for the reader who has never deeply examined a modern cryptography algorithm before.

First, algorithms like AES are typically composed of many smaller units of varying kinds. Consequently, the entire algorithm is constructed bottom-up by specifying and verifying each of its component pieces. It is wise to handle smaller and simpler components first. It is also a good practice, though hard to accomplish the first one or two times you write such a specification, to write specifications with an eye toward reuse in multiple instantiations of the same algorithm (e.g., different key sizes or configurations). The choice between encoding configurations at the type level or the value level is aesthetic and practical: some verification is only possible when one encodes information at the type level.

Second, algorithms frequently depend upon interesting data structures and mathematical constructs, the latter of which can be thought of as data structures in a pure mathematics sense. The definition, shape, form, and subtleties of these data structures are critical to the *correct definition* of the crypto algorithm *as well as its security properties*. Implementing an algorithm using an alternative datatype construction that you believe has the same properties as that which is stipulated in a standard is nearly always the wrong thing to do. Also, the subtleties of these constructions usually boils down to what an engineer might think of as “magic numbers”—strange initial values or specific polynomials that appear out of thin air. Just remind yourself that the discovery and analysis of those magic values was, in general, the joint hard work of a community of cryptographers.

6.1 Parameters

The AES algorithm always takes 128-bits of input, and always produces 128-bits of output, regardless of the key size. The key-size can be one of 128 (AES128), 192 (AES192), or 256 (AES256). Following the standard, we define the following three parameters [13, Section 2.2]:

- `Nb`: Number of columns, always set to 4 by the standard.
- `Nk`: Number of key-blocks, which is the number of 32-bit words in the key: 4 for AES128, 6 for AES192, and 8 for AES256;
- `Nr`: Number of rounds, which `Nr` is always $6 + Nk$, according to the standard. Thus, 10 for AES128, 12 for AES192, and 14 for AES256.

The Cryptol definitions follow the above descriptions verbatim:

```
type AES128 = 4
type AES192 = 6
type AES256 = 8

type Nk = AES128
type Nb = 4
type Nr = 6 + Nk
```

The following derived type is helpful in signatures:

```
type AESKeySize = (Nk*32)
```

6.2 Polynomials in $\text{GF}(2^8)$

AES works on a two-dimensional representation of the input arranged into bytes, called the *state*. For an 128-bit input, we have precisely 4 rows, each containing `Nb` (i.e., 4) bytes, each of which is 8-bits wide, totaling $4 \times 4 \times 8 = 128$ bits. The bytes themselves are treated as finite field elements in the Galois field $\text{GF}(2^8)$ [21], giving rise to the following declarations:

```
type GF28 = [8]
type State = [4][Nb]GF28
```

The hard-encoding of `GF28` in this specification is completely appropriate because the construction of AES depends entirely upon the Galois field $\text{GF}(2^8)$. It is conceivable that other algorithms might be parameterized across $\text{GF}(2^k)$ for some k , in which case the underlying type declaration would also be parameterized.

While a basic understanding Galois field operations is helpful, the details are not essential for our modeling purposes. It suffices to note that $\text{GF}(2^8)$ has precisely 256 elements, each of which is a polynomial of maximum degree 7, where the coefficients are either 0 or 1. The numbers from 0 to 255 (i.e., all possible 8-bit values) are in one-to-one correspondence with these polynomials. The coefficients of the polynomial come from the successive bits of the number, and vice versa. For instance the 8-bit number 87 can be written as `0b01010111` in binary, and hence corresponds to the polynomial $x^6 + x^4 + x^2 + x^1 + 1$. Similarly, the polynomial $x^4 + x^3 + x$ corresponds to the number `0b00011010`, i.e., 26. We can also compute this value by evaluating the polynomial for $x = 2$, obtaining $2^4 + 2^3 + 2 = 16 + 8 + 2 = 26$.

Cryptol allows you to write polynomials in $\text{GF}(2^n)$, for arbitrary n , using the following notation:

6.2. Polynomials in $GF(2^8)$

```
Cryptol> <| x^^6 + x^^4 + x^^2 + x^^1 + 1 |>
87
Cryptol> 0b1010111
87
Cryptol> <| x^^4 + x^^3 + x |>
26
Cryptol> 0b11010
26
```

A polynomial is similar to a decimal representation of a number, albeit in a more suggestive syntax. Like with a decimal, the Cryptol type system will default the type to be the smallest number of bits required to represent the polynomial, but it may be expanded to more bits if an expression requires it.

Addition and Subtraction Given two polynomials, adding and subtracting them in a Galois field $GF(2^n)$ results in a new polynomial where terms with the same power cancel each other out. When interpreted as a word, both addition and subtraction amount to a simple exclusive-or operation. Cryptol's `^` operator captures this idiom concisely:

```
Cryptol> (<| x^^4 + x^^2 |> ^ <| x^^5 + x^^2 + 1 |>) == \
<| x^^5 + x^^4 + 1 |>
True
```

Note that the term x^2 cancels since it appears in both polynomials. Also note the parentheses are required due to the precedence of `==` vs. `^`.

Exercise 1. Adding a polynomial to itself in $GF(2^n)$ will always yield 0 since all the terms will cancel each other. Write and prove a theorem `polySelfAdd` over $GF28$ to state this fact.

While adding two polynomials does not warrant a separate function, we will need a version that can add a sequence of polynomials:

Exercise 2. Define a function

```
gf28Add : {n} (fin n) => [n]GF28 -> GF28
```

that adds all the elements given. (**Hint:** Use a fold, see Pg. 29.)

Multiplication Multiplication $GF(2^n)$ follows the usual polynomial multiplication algorithm, where we multiply the first polynomial with each term of the second, and add all the partial sums (i.e., compute their exclusive-or). While this operation can be programmed explicitly, Cryptol does provide the primitive `pmult` for this purpose:

```
Cryptol> pmult <| x^^3 + x^^2 + x + 1 |> <| x^^2 + x + 1 |>
45
Cryptol> <| x^^5 + x^^3 + x^^2 + 1 |>
45
```

Exercise 3. Multiply the polynomials $x^3 + x^2 + x + 1$ and $x^2 + x + 1$ by hand in $GF(2^8)$ and show that the result is indeed $x^5 + x^3 + x^2 + 1$, (i.e., 45), justifying Cryptol's result above.

Reduction If you multiply two polynomials with degrees m and n , you will get a new polynomial of degree $m+n$. As we mentioned before, the polynomials in $\text{GF}(2^8)$ have degree at most 7. Obviously, $m+n$ can be larger than 7 when we multiply to elements of $\text{GF}(2^8)$. So we have to find a way to map the resulting larger-degree polynomial back to an element of $\text{GF}(2^8)$. This is done by reduction, or modulus, with respect to an *irreducible polynomial*. The AES algorithm uses the following polynomial for this purpose:

```
irreducible = <| x^8 + x^4 + x^3 + x + 1 |>
```

(Recall in the introduction of this chapter our warning about magic!)

Note that `irreducible` is *not* an element of $\text{GF}(2^8)$, since it has degree 8. However we can use this polynomial to define the multiplication routine itself, which uses Cryptol's `pmod` (polynomial modulus) function, as follows:

```
gf28Mult : (GF28, GF28) -> GF28
gf28Mult (x, y) = pmod (pmult x y) irreducible
```

Polynomial modulus and division operations follow the usual schoolbook algorithm for long-division—a fairly laborious process in itself, but it is well studied in mathematics [22]. Luckily for us, Cryptol's `pdiv` and `pmod` functions implement these operations, saving us the programming task.

Exercise 4. Divide $x^5 + x^3 + 1$ by $x^3 + x^2 + 1$ by hand, finding the quotient and the remainder. Check your answer with Cryptol's `pmod` and `pdiv` functions.

Exercise 5. Write and prove theorems showing that `gf28Mult` (i) has the polynomial 1 as its unit, (ii) is commutative, and (iii) is associative.

6.3 The `subBytes` transformation

Recall that the state in AES is a 4×4 matrix of bytes. As part of its operation, AES performs the so called `subBytes` transformation [13, Section 5.1.1], substituting each byte in the state with another element. Given an $x \in \text{GF}(2^8)$, the substitution for x is computed as follows:

1. Compute the multiplicative inverse of x in $\text{GF}(2^8)$, call it b . If x is 0, then take 0 as the result.
2. Replace bits in b as follows: Each bit b_i becomes $b_i \oplus b_{i+4 \pmod 8} \oplus b_{i+5 \pmod 8} \oplus b_{i+6 \pmod 8} \oplus b_{i+7 \pmod 8} \oplus c_i$. Here \oplus is exclusive-or and c is `0x63`.

Computing the multiplicative inverse It turns out that the inverse of any non-zero x in $\text{GF}(2^8)$ can be computed by raising x to the power 254, i.e., multiplying x by itself 254 times. (Mathematically, $\text{GF}(2^8)$ is a cyclic group such that x^{255} is always 1 for any x , making x^{254} the multiplicative inverse of x .)

Exercise 6. Write a function

```
gf28Pow : (GF28, [8]) -> GF28
```

such that the call `gf28Pow (n, k)` returns n^k using `gf28Mult` as the multiplication operator. (**Hint:** Use the fact that $x^0 = 1$, $x^{2n} = (x^n)^2$, and $x^{2n+1} = x \times (x^n)^2$ to speed up the exponentiation.)

Exercise 7. Write a function

6.3. The `subBytes` transformation

```
gf28Inverse : GF28 -> GF28
```

to compute the multiplicative inverse of a given element by raising it to the power 254. Note that `gf28Inverse` must map 0 to 0. Do you have to do anything special to make sure this happens?

Exercise 8. Write and prove a property `gf28InverseCorrect`, ensuring that `gf28Inverse x` does indeed return the multiplicative inverse of `x`. Do you have to do anything special when `x` is 0?

Transforming the result As we mentioned above, the AES specification asks us to transform each bit b_i according to the following transformation:

$$b_i \oplus b_{i+4 \pmod{8}} \oplus b_{i+5 \pmod{8}} \oplus b_{i+6 \pmod{8}} \oplus b_{i+7 \pmod{8}} \oplus c_i$$

For instance, bit b_5 becomes $b_5 \oplus b_1 \oplus b_2 \oplus b_3 \oplus c_5$. When interpreted at the word level, this basically amounts to computing:

$$b \oplus (b \ggg 4) \oplus (b \ggg 5) \oplus (b \ggg 6) \oplus (b \ggg 7) \oplus c$$

by aligning the corresponding bits in the word representation.

Exercise 9. Write a function

```
xformByte : GF28 -> GF28
```

that computes the above described transformation.

Putting it together Armed with `gf28Inverse` and `xformByte`, we can easily code the function that transforms a single byte as follows:

```
SubByte : GF28 -> GF28
SubByte b = xformByte (gf28Inverse b)
```

AES's `subBytes` transformation merely applies this function to each byte of the state:

```
SubBytes : State -> State
SubBytes state = [ [ SubByte b | b <- row ] | row <- state ]
```

Table lookup Our definition of the `subByte` function above follows how the designers of AES came up with the substitution maps, i.e., it is a *reference* implementation. For efficiency purposes, however, we might prefer a version that simply performs a look-up in a table. Notice that the type of `SubByte` is `GF28 -> GF28`, i.e., it takes a value between 0 and 255. Therefore, we can make a table containing the precomputed values for all possible 256 inputs, and simply perform a table look-up instead of computing these values each time we need it. In fact, Figure 7 on page 16 of the AES standard lists these precomputed values for us [13, Section

6.4. The `shiftRows` transformation

```
sbox : [256]GF28
sbox = [
  0x63, 0x7c, 0x77, 0x7b, 0xf2, 0x6b, 0x6f, 0xc5, 0x30, 0x01, 0x67,
  0x2b, 0xfe, 0xd7, 0xab, 0x76, 0xca, 0x82, 0xc9, 0x7d, 0xfa, 0x59,
  0x47, 0xf0, 0xad, 0xd4, 0xa2, 0xaf, 0x9c, 0xa4, 0x72, 0xc0, 0xb7,
  0xfd, 0x93, 0x26, 0x36, 0x3f, 0xf7, 0xcc, 0x34, 0xa5, 0xe5, 0xf1,
  0x71, 0xd8, 0x31, 0x15, 0x04, 0xc7, 0x23, 0xc3, 0x18, 0x96, 0x05,
  0x9a, 0x07, 0x12, 0x80, 0xe2, 0xeb, 0x27, 0xb2, 0x75, 0x09, 0x83,
  0x2c, 0x1a, 0x1b, 0x6e, 0x5a, 0xa0, 0x52, 0x3b, 0xd6, 0xb3, 0x29,
  0xe3, 0x2f, 0x84, 0x53, 0xd1, 0x00, 0xed, 0x20, 0xfc, 0xb1, 0x5b,
  0x6a, 0xcb, 0xbe, 0x39, 0x4a, 0x4c, 0x58, 0xcf, 0xd0, 0xef, 0xaa,
  0xfb, 0x43, 0x4d, 0x33, 0x85, 0x45, 0xf9, 0x02, 0x7f, 0x50, 0x3c,
  0x9f, 0xa8, 0x51, 0xa3, 0x40, 0x8f, 0x92, 0x9d, 0x38, 0xf5, 0xbc,
  0xb6, 0xda, 0x21, 0x10, 0xff, 0xf3, 0xd2, 0xcd, 0x0c, 0x13, 0xec,
  0x5f, 0x97, 0x44, 0x17, 0xc4, 0xa7, 0x7e, 0x3d, 0x64, 0x5d, 0x19,
  0x73, 0x60, 0x81, 0x4f, 0xdc, 0x22, 0x2a, 0x90, 0x88, 0x46, 0xee,
  0xb8, 0x14, 0xde, 0x5e, 0x0b, 0xdb, 0xe0, 0x32, 0x3a, 0x0a, 0x49,
  0x06, 0x24, 0x5c, 0xc2, 0xd3, 0xac, 0x62, 0x91, 0x95, 0xe4, 0x79,
  0xe7, 0xc8, 0x37, 0x6d, 0x8d, 0xd5, 0x4e, 0xa9, 0x6c, 0x56, 0xf4,
  0xea, 0x65, 0x7a, 0xae, 0x08, 0xba, 0x78, 0x25, 0x2e, 0x1c, 0xa6,
  0xb4, 0xc6, 0xe8, 0xdd, 0x74, 0x1f, 0x4b, 0xbd, 0x8b, 0x8a, 0x70,
  0x3e, 0xb5, 0x66, 0x48, 0x03, 0xf6, 0x0e, 0x61, 0x35, 0x57, 0xb9,
  0x86, 0xc1, 0x1d, 0x9e, 0xe1, 0xf8, 0x98, 0x11, 0x69, 0xd9, 0x8e,
  0x94, 0x9b, 0x1e, 0x87, 0xe9, 0xce, 0x55, 0x28, 0xdf, 0x8c, 0xa1,
  0x89, 0x0d, 0xbf, 0xe6, 0x42, 0x68, 0x41, 0x99, 0x2d, 0x0f, 0xb0,
  0x54, 0xbb, 0x16]
```

5.1.1]. We capture this table below in Cryptol:

With this definition of `sbox`, the look up variants of `SubByte` and `SubBytes` becomes:

```
SubByte' : GF28 -> GF28
SubByte' x = sbox @ x

SubBytes' : State -> State
SubBytes' state = [ [ SubByte' b | b <- row ] | row <- state ]
```

Exercise 10. Write and prove a property stating that `SubByte` and `SubByte'` are equivalent.

NB. The `SubByte'` and `SubBytes'` versions are going to be more efficient for execution, naturally. We should emphasize that this mode of development is quite common in modern cryptography. Ciphers are typically designed using ideas from mathematics, often requiring complicated algorithms. To speed things up, however, implementors use clever optimization tricks, convert functions to look-up tables using precomputed values, etc.

What Cryptol allows us to do is to write the algorithms using both styles, and then formally show that they are indeed equivalent, as you did in Exercise 10 above. This mode of high-assurance development makes sure that we have not made any cut-and-paste errors when we wrote down the numbers in our `sbox` table. Equivalently, our proof also establishes that the official specification [13, Section 5.1.1] got its own table correct!

6.4 The `shiftRows` transformation

The second transformation AES utilizes is the `shiftRows` operation [13, Section 5.1.2]. This operation treats the `state` as a 4×4 matrix, and rotates the last three row to the left by the amounts 1, 2, and 3; respectively. Implementing `shiftRows` in Cryptol is trivial, using the `<<<` operator:

6.5. The `mixColumns` transformation

```
ShiftRows : State -> State
ShiftRows state = [ row <<< shiftAmount | row <- state
                  | shiftAmount <- [0 .. 3]
                  ]
```

Exercise 11. Can you transform a state back into itself by repeated applications of `ShiftRows`? How many times would you need to shift? Verify your answer by writing and proving a corresponding Cryptol property.

6.5 The `mixColumns` transformation

The third major transformation AES performs is the `mixColumns` operation [13, Section 5.1.3]. In this transformation, the `State` is viewed as a 4×4 matrix, and each successive column of it is replaced by the result of multiplying it by the matrix:

$$\begin{bmatrix} 2 & 3 & 1 & 1 \\ 1 & 2 & 3 & 1 \\ 1 & 1 & 2 & 3 \\ 3 & 1 & 1 & 2 \end{bmatrix}$$

As you might recall from linear algebra, given two *compatible* matrices A and B , the ij 'th element of $A \times B$ is the dot-product of the i 'th row of A and the j 'th column of B . (By *compatible* we mean the number of columns of A must be the same as the number of rows of B . All our matrices are 4×4 , so they are always compatible.) The dot-product is defined as multiplying the corresponding elements of two same-length vectors and adding the results together. The only difference here is that we use the functions `gf28Add` and `gf28Mult` for addition and multiplication respectively. We will develop this algorithm in the following sequence of exercises.

Exercise 12. Write a function `gf28DotProduct` with the signature:

```
gf28DotProduct : {n} (fin n) => ([n]GF28, [n]GF28) -> GF28
```

such that `gf28DotProduct` returns the dot-product of two length n vectors of $\text{GF}(2^8)$ elements.

Exercise 13. Write properties stating that the dot-matrix operation `gf28DotProduct` is commutative and distributive over vector addition:

$$\begin{aligned} a \cdot b &= b \cdot a \\ a \cdot (b + c) &= a \cdot b + a \cdot c \end{aligned}$$

Addition over vectors is defined element-wise. Prove the commutativity property for vectors of length 10. Distributivity will take much longer, so you might want to do a `:check` on it.

Exercise 14. Write a function

```
gf28VectorMult : {n, m, k} (fin n) => ([n]GF28, [m][n]GF28) ->
                                     [m]GF28
```

computing the dot-product of its first argument with each of the m rows of the second argument, returning the resulting values as a vector of m elements.

Exercise 15. Write a function

```
gf28MatrixMult : {n, m, k} (fin m) => ([n][m]GF28, [m][k]GF28)
               -> [n][k]GF28
```

which multiplies the given matrices in $\text{GF}(2^8)$.

Now that we have the matrix multiplication machinery built, we can code `MixColumns` fairly easily. Following the description in the AES standard [13, Section 5.3.1], all we have to do is to multiply the matrix we have seen at the beginning of this section with the state:

```
MixColumns : State -> State
MixColumns state = gf28MatrixMult (m, state)
  where m = [[2, 3, 1, 1],
             [1, 2, 3, 1],
             [1, 1, 2, 3],
             [3, 1, 1, 2]]
```

Note that Cryptol makes no built-in assumption about row- or column-ordering of multidimensional matrices. Of course, given Cryptol's concrete syntax, it makes little sense to do anything but row-based ordering.

6.6 Key expansion

Recall from Section 6.1 that AES takes 128, 192, or 256 bit keys. The key is not used as-is, however. Instead, AES expands the key into a number of round keys, called the *key schedule*. Construction of the key schedule relies on a number of auxiliary definitions, as we shall see shortly.

Round constants AES standard refers to the constant array `Rcon` used during key expansion. For each i , `Rcon[i]` contains 4 words, the last three being 0 [13, Section 5.2]. The first element is given by x^{i-1} , where exponentiation is done using the `gf28Pow` function you have defined in Exercise 6.3-6. In Cryptol, it is easiest to define `Rcon` as a function:

```
Rcon : [8] -> [4]GF28
Rcon i = [(gf28Pow (<| x |>, i-1)), 0, 0, 0]
```

Exercise 16. By definition, AES only calls `Rcon` with the parameters ranging from 1–10. Based on this, create a table-lookup version

```
Rcon' : [8] -> [4]GF28
```

that simply performs a look-up instead. (**Hint:** Use Cryptol to find out what the elements of your table should be.)

Exercise 17. Write and prove a property that `Rcon` and `Rcon'` are equivalent when called with numbers in the range 1–10.

6.6. Key expansion

The subword function AES specification refers to a function named `SubWord` [13, Section 5.2], that takes a 32-bit word and applies the `SubByte` transformation from Section 6.3. This function is trivial to code in Cryptol:

```
SubWord : [4]GF28 -> [4]GF28
SubWord bs = [ SubByte' b | b <- bs ]
```

Note that we have used the table-lookup version (`SubByte'`, Pg 66) above.

The rotword function The last function we need for key-expansion is named `RotWord` by the AES standard [13, Section 5.2]. This function merely rotates a given word cyclically to the left:

```
RotWord : [4]GF28 -> [4]GF28
RotWord [a0, a1, a2, a3] = [a1, a2, a3, a0]
```

We could have used `<<<` to implement `RotWord` as well, but the above definition textually looks exactly the one given in the standard specification, and hence is preferable for the purposes of clarity.

The key schedule Recall that AES operates on 128, 192, or 256 bit keys. These keys are used to construct a sequence of so called *round-keys*, each of which is 128-bits wide, and is viewed the same way as the `State`:

```
type RoundKey = State
```

The expanded key schedule, contains N_r+1 round-keys. (Recall from Section 6.1 that N_r is the number of rounds.) It also helps to separate out the first and the last keys, as they are used in a slightly different fashion. Based on this discussion, we use the following Cryptol type to capture the key-schedule:

```
type KeySchedule = (RoundKey, [Nr-1]RoundKey, RoundKey)
```

The key schedule is computed by seeding it with the initial key and computing the successive elements from the previous entries. In particular, the i th element of the expanded key is determined as follows, copied verbatim from the AES specification [13, Figure 11; Section 5.2]:

```
temp = w[i-1]
if (i mod Nk = 0)
  temp = SubWord(RotWord(temp)) xor Rcon[i/Nk]
else if (Nk > 6 and i mod Nk = 4)
  temp = SubWord(temp)
end if
w[i] = w[i-Nk] xor temp
```

In the pseudo-code, the `w` array contains the expanded key. We are computing `w[i]`, using the values `w[i-1]` and `w[i-Nk]`. The result is the exclusive-or of `w[i-Nk]` and a mask value, called `temp` above. The mask is computed using `w[i-1]`, the `Rcon` array we have seen before, the current index `i`, and `Nk`. This computation is best expressed as a function in Cryptol that we will call `NextWord`. We will name the `w[i-1]` argument `prev`, and the `w[i-Nk]` argument `old`. Otherwise, the Cryptol code just follows the pseudo-code above, written in a functional style to compute the mask:

```

NextWord : ([8],[4][8],[4][8]) -> [4][8]
NextWord(i, prev, old) = old ^ mask
  where mask = if i % 'Nk == 0
                then SubWord(RotWord(prev)) ^ Rcon' (i / 'Nk)
                else if ('Nk > 6) && (i % 'Nk == 4)
                then SubWord(prev)
                else prev

```

NB. It is well worth studying the pseudo-code above and the Cryptol equivalent to convince yourself they are expressing the same idea!

To compute the key schedule we start with the initial key as the seed. We then make calls to `NextWord` with a sliding window of `Nk` elements, computing the subsequent elements. Let us first write a function that will apply this algorithm to generate an infinite regression of elements:

```

ExpandKeyForever : [Nk][4][8] -> [inf]RoundKey
ExpandKeyForever seed = [ transpose g | g <- groupBy' {4} keyWS ]
  where keyWS = seed # [ NextWord(i, prev, old)
                        | i <- ['Nk ...]
                        | prev <- drop' {Nk-1} keyWS
                        | old <- keyWS
                    ]

```

Note how `prev` tracks the previous 32-bits of the expanded key (by dropping the first `Nk-1` elements), while `old` tracks the `i-Nk`'th recurrence for `keyWS`. Once we have the infinite expansion, it is easy to extract just the amount we need by using number of rounds (`Nr`) as our guide:

```

ExpandKey : [AESKeySize] -> KeySchedule
ExpandKey key = (keys @ 0, keys @@ [1 .. (Nr - 1)], keys @ 'Nr)
  where seed : [Nk][4][8]
        seed = split (split key)
        keys = ExpandKeyForever seed

```

The call `split key` chops `AESKey` into `[Nk*4][8]`, and the outer call to `split` further constructs the `[Nk][4][8]` elements.

Testing `ExpandKey` The completion of `ExpandKey` is an important milestone in our AES development, and it is worth testing it before we proceed. The AES specification has example key-expansions that we can use. The following function will be handy in viewing the output correctly aligned:

```

fromKS : KeySchedule -> [Nr+1][4][32]
fromKS (f, ms, l) = [ formKeyWords (transpose k)
                    | k <- [f] # ms # [l]
                    ]
  where formKeyWords bbs = [ join bs | bs <- bbs ]

```

Here is the example from Appendix A.1 of the AES specification [13]:

```

Cryptol> fromKS (ExpandKey 0x2b7e151628aed2a6abf7158809cf4f3c)
[[0x2b7e1516, 0x28aed2a6, 0xabf71588, 0x09cf4f3c],

```

6.7. The `AddRoundKey` transformation

```
[0xa0fafa17, 0x88542cb1, 0x23a33939, 0x2a6c7605],
[0xf2c295f2, 0x7a96b943, 0x5935807a, 0x7359f67f],
[0x3d80477d, 0x4716fe3e, 0x1e237e44, 0x6d7a883b],
[0xef44a541, 0xa8525b7f, 0xb671253b, 0xdb0bad00],
[0xd4d1c6f8, 0x7c839d87, 0xcaf2b8bc, 0x11f915bc],
[0x6d88a37a, 0x110b3efd, 0xdbf98641, 0xca0093fd],
[0x4e54f70e, 0x5f5fc9f3, 0x84a64fb2, 0x4ea6dc4f],
[0xead27321, 0xb58dbad2, 0x312bf560, 0x7f8d292f],
[0xac7766f3, 0x19fadc21, 0x28d12941, 0x575c006e],
[0xd014f9a8, 0xc9ee2589, 0xe13f0cc8, 0xb6630ca6]
```

As you can verify this output matches the last column of the table in Appendix A.1 of the reference specification for AES.

6.7 The `AddRoundKey` transformation

`AddRoundKey` is the simplest of all the transformations in AES [13, Section 5.1.4]. It merely amounts to the exclusive-or of the state and the current round key:

```
AddRoundKey : (RoundKey, State) -> State
AddRoundKey (rk, s) = rk ^ s
```

Notice that Cryptol's `^` operator applies structurally to arbitrary shapes, computing the exclusive-or element-wise.

6.8 AES encryption

We now have all the necessary machinery to perform AES encryption.

AES rounds As mentioned before, AES performs encryption in rounds. Each round consists of performing `SubBytes` (Section 6.3), `ShiftRows` (Section 6.4), and `MixColumns` (Section 6.5). Before finishing up, each round also adds the current round key to the state [13, Section 5.1]. The Cryptol code for the rounds is fairly trivial:

```
AESRound : (RoundKey, State) -> State
AESRound (rk, s) = AddRoundKey (rk,
                               MixColumns (ShiftRows (SubBytes s)))
```

The final round The last round of AES is slightly different than the others. It omits the `MixColumns` transformation:

```
AESFinalRound : (RoundKey, State) -> State
AESFinalRound (rk, s) = AddRoundKey (rk, ShiftRows (SubBytes s))
```

Forming the input/output blocks Recall that AES processes input in blocks of 128-bits, producing 128-bits of output, regardless of the key size. We will need two helper functions to convert 128-bit messages to and from AES states. Conversion from a message to a state is easy to define:

```
msgToState : [128] -> State
msgToState msg = transpose (split (split msg))
```

The first call to `split` gives us 4 32-bit words, which we again split into bytes. We then form the AES state by transposing the resulting matrix. In the other direction, we simply transpose the state and perform the necessary join's:

```
stateToMsg : State -> [128]
stateToMsg st = join (join (transpose st))
```

Exercise 18. Write and prove a pair of properties stating that `msgToState` and `stateToMsg` are inverses of each other.

Putting it together To encrypt, AES merely expands the given key and calls the round-functions. The starting state (`state0` below) is constructed by adding the first round key to the input. We then run all the middle rounds using a simple comprehension, and finish up by applying the last round [13, Figure 5, Section 5.1]:

```
aesEncrypt : ([128], [AESKeySize]) -> [128]
aesEncrypt (pt, key) = stateToMsg (AESFinalRound (kFinal, rounds ! 0))
  where (kInit, ks, kFinal) = ExpandKey key
        state0 = AddRoundKey(kInit, msgToState pt)
        rounds = [state0] # [ AESRound (rk, s) | rk <- ks
                               | s <- rounds
                               ]
```

Testing We can now run some test vectors. Note that, just because a handful of test vectors pass, we cannot claim that our implementation of AES is correct.

The first example comes from Appendix B of the AES standard: [13]:

```
Cryptol> aesEncrypt (0x3243f6a8885a308d313198a2e0370734, \
                    0x2b7e151628aed2a6abf7158809cf4f3c)
0x3925841d02dc09fbdc118597196a0b32
```

which is what the standard asserts to be the answer. (Note that you have to read the final box in Appendix B column-wise!) The second example comes from Appendix C.1:

```
Cryptol> aesEncrypt (0x00112233445566778899aabbccddeeff, \
                    0x000102030405060708090a0b0c0d0e0f)
0x69c4e0d86a7b0430d8cdb78070b4c55a
```

Again, the result agrees with the standard.

6.9. Decryption

Other key sizes Our development of AES has been key-size agnostic, relying on the definition of the parameter n_k (See Section 6.1.) To obtain AES192, all we need is to set n_k to be 6, no additional code change is needed. Similarly, we merely need to set n_k to be 8 for AES256.

Exercise 19. By setting n_k to be 6 and 8 respectively, try the test vectors given in Appendices C.2 and C.3 of the AES standard [13].

6.9 Decryption

AES decryption is fairly similar to encryption, except it uses inverse transformations [13, Figure 12, Section 5.3]. Armed with all the machinery we have built so far, the inverse transformations is relatively easy to define.

6.9.1 The `InvSubBytes` transformation

The `InvSubBytes` transformation reverses the `SubBytes` transformation of Section 6.3. As with `SubBytes`, we have a choice to either do a table lookup implementation, or follow the mathematical description. We will do the former in these examples; you are welcome to do the latter on your own and prove the equivalence of the two versions. To do so, we need to invert the transformation given by:

$$b \oplus b \ggg 4 \oplus b \ggg 5 \oplus b \ggg 6 \oplus b \ggg 7 \oplus c$$

where c is `0x63`. It turns out that the inverse of this transformation can be computed by

$$b \ggg 2 \oplus b \ggg 5 \oplus b \ggg 7 \oplus d$$

where d is `0x05`. It is easy to code this inverse transform in Cryptol:

```
xformByte' : GF28 -> GF28
xformByte' b = gf28Add [(b >>> 2), (b >>> 5), (b >>> 7), d]
  where d = 0x05
```

Exercise 20. Write and prove a Cryptol property stating that `xformByte'` is the inverse of the function `xformByte` that you have defined in Exercise 6.3 9.

We can now define the inverse s-box transform, using the multiplicative inverse function `gf28Inverse` you have defined in Exercise 6.3 7:

```
InvSubByte : GF28 -> GF28
InvSubByte b = gf28Inverse (xformByte' b)

InvSubBytes : State -> State
InvSubBytes state = [ [ InvSubByte b | b <- row ]
                      | row <- state
                      ]
```

Exercise 21. Write and prove a Cryptol property showing that `InvSubByte` reverses `SubByte`.

Exercise 22. The AES specification provides an inverse s-box table [13, Figure 14, Section 5.3.2]. Write a Cryptol function `InvSubBytes'` using the table lookup technique. Make sure your implementation is correct (i.e., equivalent to `InvSubBytes`) by writing and proving a corresponding property.

6.9.2 The `InvShiftRows` transformation

The `InvShiftRows` transformation simply reverses the `ShiftRows` transformation from Section 6.4:

```
InvShiftRows : State -> State
InvShiftRows state = [ row >>> shiftAmount
                      | row <- state
                      | shiftAmount <- [0 .. 3]
                      ]
```

Exercise 23. Write and prove a property stating that `InvShiftRows` is the inverse of `ShiftRows`.

6.9.3 The `InvMixColumns` transformation

Recall from Section 6.5 that `MixColumns` amounts to matrix multiplication in $GF(2^8)$. The inverse transform turns out to be the same, except with a different matrix:

```
InvMixColumns : State -> State
InvMixColumns state = gf28MatrixMult (m, state)
  where m = [[0x0e, 0x0b, 0x0d, 0x09],
             [0x09, 0x0e, 0x0b, 0x0d],
             [0x0d, 0x09, 0x0e, 0x0b],
             [0x0b, 0x0d, 0x09, 0x0e]]
```

Exercise 24. Write and prove a property stating that `InvMixColumns` is the inverse of `MixColumns`.

6.10 The inverse cipher

We now also have all the ingredients to encode AES decryption. Following Figure 12 (Section 5.3) of the AES standard [13]:

```
AESInvRound : (RoundKey, State) -> State
AESInvRound (rk, s) =
  InvMixColumns (AddRoundKey (rk, InvSubBytes (InvShiftRows s)))

AESFinalInvRound : (RoundKey, State) -> State
AESFinalInvRound (rk, s) = AddRoundKey (rk, InvSubBytes (InvShiftRows s))

aesDecrypt : ([128], [AESKeySize]) -> [128]
aesDecrypt (ct, key) = stateToMsg (AESFinalInvRound (kFinal, rounds ! 0))
  where
```

6.11. Correctness

```
(kFinal, ks, kInit) = ExpandKey key
state0 = AddRoundKey(kInit, msgToState ct)
rounds = [state0] # [ AESInvRound (rk, s)
                  | rk <- reverse ks
                  | s <- rounds
                  ]
```

Note how we use the results of `ExpandedKey`, by carefully naming the first and last round keys and using the middle-keys in reverse.

Testing Let us repeat the tests for AES encryption. Again, the first example comes from Appendix B of the AES standard [13]:

```
Cryptol> aesDecrypt (0x3925841d02dc09fbdc118597196a0b32, \
                    0x2b7e151628aed2a6abf7158809cf4f3c)
0x3243f6a8885a308d313198a2e0370734
```

which agrees with the original value. The second example comes from Appendix C.1:

```
Cryptol> aesDecrypt (0x69c4e0d86a7b0430d8cdb78070b4c55a, \
                    0x000102030405060708090a0b0c0d0e0f)
0x00112233445566778899aabbccddeeff
```

Again, the result agrees with the standard.

Other key sizes Similar to encryption, all we need to obtain AES192 decryption is to set `nk` to be 6 in Section 6.1. Setting `nk` to 8 will correspondingly give us AES256.

The code You can see all the Cryptol code for AES in Appendix D.

6.11 Correctness

While test vectors do provide good evidence of AES working correctly, they do not provide a proof that we have implemented the standard faithfully. In fact, for a block-cipher like AES, it is not possible to state what correctness would mean. Tweaking some parameters, or changing the s-box appropriately can give us a brand new cipher. And it would be impossible to tell this new cipher apart from AES aside from running it against published test vectors.

What we can do, however, is gain assurance that our implementation demonstrably has the desired properties. We have done this throughout this chapter by stating and proving a number of properties about AES and its constituent parts. The Cryptol methodology allows us to construct the code together with expected properties, allowing high-assurance development. We conclude this chapter with one final property, stating that `aesEncrypt` and `aesDecrypt` do indeed form an encryption-decryption pair:

```
property AESCorrect msg key = aesDecrypt (aesEncrypt (msg, key), key) == msg
```

Can we hope to automatically prove this theorem? For 128-bit AES, the state space for this theorem has 2^{256} elements. It would be naive to expect that we can prove this theorem by a push-button tool very quickly.¹ We can however, gain some assurance by running it through the `:check` command:

¹Note that, for a general algorithm with this large a state space, it is entirely possible to perform automatic verification using modern solvers, but if one carefully reflects upon the nature of cryptographic functions, it becomes clear why it should *not* be the case here.

```
Checking case 1000 of 1000 (100.00%)  
1000 tests passed OK
```

You will notice that even running quick-check will take a while for the above theorem, and the total state space for this function means that we have not even scratched the surface! That said, being able to specify these properties together with very high level code is what distinguishes Cryptol from other languages when it comes to cryptographic algorithm programming.

Appendix A

Solutions to selected exercises

As with any language, there are usually multiple ways to write the same function in Cryptol. We have tried to use the most idiomatic Cryptol code segments in our solutions. Note that Cryptol prints numbers out in hexadecimal by default. In most of the answers below, we have implicitly used the command `:set base=10` to print numbers out in decimal for readability.

Section 1.1 Getting started (p.3)

Exercise 1. (p. 4) Here is the response I get:

```
Cryptol> "Hello World!"  
[0x48, 0x65, 0x6c, 0x6c, 0x6f, 0x20, 0x57, 0x6f, 0x72, 0x6c, 0x64, 0x21]
```

As we shall see in Section 2.7, strings are just sequences of ASCII characters, so Cryptol is telling you the ASCII numbers corresponding to the letters.

Exercise 2. (p. 4) In this case we see:

```
Cryptol> :set ascii=on  
Cryptol> "Hello World!"  
"Hello World!"
```

The command `:set ascii=on` asks Cryptol to treat sequences of 8 bit values as strings, which is how strings are really represented in Cryptol. This is not the default behavior, however, since sequences of 8 bit values can represent other things, especially in the domain of cryptography. The first behavior is typically what a crypto-programmer wants to see, and hence is the default.

Section 2.2 Bits: Booleans (p.9)

Exercise 1. (p. 9) Here is the response from Cryptol, in order:

```
True
[error] at <interactive>:1:1--1:8: Variable 'false' was not defined.
False
0x4
True
True
True
True
False
```

Section 2.3 Words: Numbers (p.10)

Exercise 2. (p. 10) `0xfeedfacef00d`. Allowing base 1 would have resulted in unreadable output, and anything larger than 36 would have required Cryptol to use unicode for digits (and would have limited utility). As usual a is 10, b is 11, ... z is 36. Remember that upper and lower case letters denote the same value, so a and A both represent 10.

Section 2.4 Tuples: Heterogeneous collections (p.11)

Exercise 3. (p. 11) Here are Cryptol's responses:

```
(1, 6)
(True, False, True)
((1, 2), False, (2, (4, True)))
```

Exercise 4. (p. 11) Here are Cryptol's responses:

```
Cryptol> (1, 2+4).0
1
Cryptol> (1, 2+4).1
6
Cryptol> ((1, 2), False, (3-1, (4, True))).2
(2, (4, True))
```

The required expression would be:

```
((1, 2), (2, (4, True), 6), False).2
```

Section 2.5 Sequences: Homogeneous collections (p.11)

Exercise 6. (p. 12) In each case we get a type-error:

```

Cryptol> [1, True]
[error] at <interactive>:1:1--1:10:
  Type mismatch:
    Expected type: [?a]
    Inferred type: Bit
Cryptol> [[1, 2, 3], [4, 5]]
[error] at <interactive>:1:1--1:20:
  Type mismatch:
    Expected type: [3][?a]
    Inferred type: [2][?b]

```

In the first case, we are trying to put a bit (`True`) and a singleton sequence containing a bit (`[True]`) in the same sequence, which have different types. In the second case, we are trying to put two sequences of different lengths into a sequence, which again breaks the homogeneity requirement.

Exercise 7. (p. 12) Here are the responses from Cryptol:

```

[1, 2, 3, 4, 5, 6, 7, 8, 9, 10]
[1, 3, 5, 7, 9]
[10, 9, 8, 7, 6, 5, 4, 3, 2, 1]
[]
[10, 7, 4, 1]
[]

```

Note how `[10, 11 .. 1]` and `[10, 9 .. 20]` give us empty sequences, since the upper bound is smaller than the lower bound in the former, and larger in the latter.

Exercise 8. (p. 12) Here are the responses from Cryptol:

```

[(1, 4) (1, 5) (2, 4) (2, 5) (3, 4) (3, 5)]
[]
[(2, 3) (2, 4) (3, 3) (3, 4)]

```

The size of the result will be the sizes of the components multiplied. For instance, in the first example, the generator `x <- [1 .. 3]` assigns 3 values to `x`, and the generator `y <- [4, 5]` assigns 2 values to `y`; and hence the result has $2 \times 3 = 6$ elements.

Exercise 9. (p. 13) Here are the responses from Cryptol:

```

[(1, 4) (2, 5)]
[]
[(2, 3)]

```

In this case, the size of the result will be the minimum of the component sizes. For the first example, the generator `x <- [1 .. 3]` assigns 3 values to `x`, and the generator `y <- [4, 5]` assigns 2 values to `y`; and hence the result has $\min(2, 3) = 2$ elements.

Exercise 10. (p. 13) Here is one way of writing such an expression, layed out in multiple lines to show the structure:

```
[ [ (i, j) | j <- [1 .. 3] ]
  | i <- [1 .. 3]
]
produces:
[[ (1, 1), (1, 2), (1, 3) ],
 [ (2, 1), (2, 2), (2, 3) ],
 [ (3, 1), (3, 2), (3, 3) ]]
```

The outer comprehension is a comprehension (and hence is nested). In particular the expression is:

```
[ (i, j) | j <- [1 .. 3] ]
```

You can enter the whole expression in Cryptol all in one line, or recall that you can put `\` at line ends to continue to the next line. If you are writing such an expression in a program file, then you can lay it out as shown above or however most makes sense to you.

Exercise 11. (p. 13) Here are Cryptol's responses:

```
Cryptol> :set warnDefaulting=off
Cryptol> [] # [1, 2]
[1, 2]
Cryptol> [1, 2] # []
[1, 2]
Cryptol> [1 .. 5] # [3, 6, 8]
[1, 2, 3, 4, 5, 3, 6, 8]
Cryptol> [0 .. 9] @ 0
0
Cryptol> [0 .. 9] @ 5
5
Cryptol> [0 .. 9] @ 10
invalid sequence index: 10
Cryptol> [0 .. 9] @@ [3, 4]
[3, 4]
Cryptol> [0 .. 9] @@ []
[]
Cryptol> [0 .. 9] @@ [9, 12]
invalid sequence index: 12
Cryptol> [0 .. 9] @@ [9, 8 .. 0]
[9, 8, 7, 6, 5, 4, 3, 2, 1, 0]
Cryptol> [0 .. 9] ! 0
9
Cryptol> [0 .. 9] ! 3
6
Cryptol> [0 .. 9] !! [3, 6]
[6, 3]
Cryptol> [0 .. 9] !! [0 .. 9]
[9, 8, 7, 6, 5, 4, 3, 2, 1, 0]
Cryptol> [0 .. 9] ! 12
invalid sequence index: 12
```


Exercise 12. (p. 14) Using a permutation operator, we can simply write:

```
[0 .. 10] @@ [0, 2 .. 10]
```

Using a comprehension, we can express the same idea using:

```
[ [0 .. 10] @ i | i <- [0, 2 .. 10] ]
```

Strictly speaking, permutation operations are indeed redundant. However, they lead to more concise and easier-to-read expressions.

Exercise 13. (p. 14) When you type in an infinite sequence, Cryptol will only print the first 5 elements of it and will indicate that it is an infinite value by putting ... at the end³. Here are the responses:

```
[1, 2, 3, 4, 5, ...]
[1, 3, 5, 7, 9, ...]
2001
[601, 1001, 1401]
[100, 102, 104, 106, 108, ...]
```

Exercise 14. (p. 14) Here is a simple test case:

```
Cryptol> ([1 ... ]:[inf])[32]!3
[error] at <interactive>:1:1--1:21:
  Unsolved constraint:
    fin inf
    arising from
    use of expression (!)
    at <interactive>:1:19--1:20
```

The error message is telling us that we *cannot* apply the reverse index operator (!) on an infinite sequence (inf). This is a natural consequence of the fact that one can never reach the end of an infinite sequence to count backwards. It is important to emphasize that this is a *type-error*, i.e., the user gets this message at compile time; instead of Cryptol going into an infinite loop to reach the end of an infinite sequence.

Exercise 15. (p. 14) Here are Cryptol's responses:

```
Cryptol> take' {3} [1 .. 12]
[1, 2, 3]
Cryptol> drop' {3} [1 .. 12]
[4, 5, 6, 7, 8, 9, 10, 11, 12]
Cryptol> split' {3} [1 .. 12]
[[1, 2, 3, 4], [5, 6, 7, 8], [9, 10, 11, 12]]
Cryptol> groupBy' {3} [1 .. 12]
[[1, 2, 3], [4, 5, 6], [7, 8, 9], [10, 11, 12]]
Cryptol> join [[1 .. 4], [5 .. 8], [9 .. 12]]
```

³You can change this behavior by setting the `infLength` variable, like so: `Cryptol> :set infLength=10` will show the first 10 elements of infinite sequences

```

[1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12]
Cryptol> join [[1, 2, 3], [4, 5, 6], [7, 8, 9], [10, 11, 12]]
[1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12]
Cryptol> transpose [[1, 2, 3, 4], [5, 6, 7, 8]]
[[1, 5], [2, 6], [3, 7], [4, 8]]
Cryptol> transpose [[1, 2, 3], [4, 5, 6], [7, 8, 9]]
[[1, 4, 7], [2, 5, 8], [3, 6, 9]]

```

Exercise 16. (p. 15) The following equalities are the simplest candidates:

```

join (split' {parts=n} xs) == xs
join (groupBy' {each=n} xs) == xs
split' {parts=n} xs == groupBy' {each=m} xs
transpose (transpose xs) == xs

```

In the first two equalities n must be a divisor of the length of the sequence xs . In the third equation, $n \times m$ must equal the length of the sequence xs .

Exercise 17. (p. 15) Append (`#`) joins two sequences of arbitrary length, while `join` appends a sequence of equal length sequences. In particular, the equality:

```

join [xs0, xs1, xs2, .. xsN] = xs0 # xs1 # xs2 ... # xsN

```

holds for all equal length sequences xs_0, xs_1, \dots, xs_N .

Exercise 18. (p. 15) Here they are:

```

Cryptol> split [1..12] : [4][3][8]
[[1, 2, 3], [4, 5, 6], [7, 8, 9], [10, 11, 12]]
Cryptol> split [1..12] : [6][2][8]
[[1, 2], [3, 4], [5, 6], [7, 8], [9, 10], [11, 12]]
Cryptol> split [1..12] : [12][1][8]
[[1], [2], [3], [4], [5], [6], [7], [8], [9], [10], [11], [12]]

```

Exercise 19. (p. 15) Cryptol will issue a type error:

```

Cryptol> split [1..12] : [5][2][8]
Unsolved constraint:
  1 + (12 - 1) == 5 * 2
  arising from
  matching types
  at <interactive>:1:1--1:16

```

Cryptol is telling us that we have requested 10 elements in the final result (5×2), but the input has 12.

Exercise 20. (p. 15) We can split 120 elements first into 3–40, splitting each of the the elements (`level1` below) into 4–10. A nested comprehension fits the bill:

```

[ split level1 : [4][10][8]
| level1 <- split ([1 .. 120] : [120][8]) : [3][40][8]
]

```

(Note again that you can enter the above in the command line all in one line, or by putting the line continuation character `\` at the end of the first two lines.)

Exercise 21. (p. 16) Here are Cryptol's responses:

```
[0, 0, 1, 2, 3]
[0, 0, 0, 0, 0]
[3, 4, 5, 0, 0]
[0, 0, 0, 0, 0]
[4, 5, 1, 2, 3]
[1, 2, 3, 4, 5]
[3, 4, 5, 1, 2]
[1, 2, 3, 4, 5]
```

Exercise 22. (p. 16) Rotating (left or right) by a multiple of the size of a sequence will leave it unchanged.

Section 2.6 Words revisited (p.16)

Exercise 23. (p. 16) Cryptol is big-endian, meaning that the most-significant-bit comes first. In the sequence `[True, False, True, False, True, False]`, the first element corresponds to the most-significant-digit, i.e., 2^5 , the next element corresponds to the coefficient of 2^4 , etc. A `False` bit yields a coefficient of 0 and a `True` bit gives 1. Hence, we have:

$$1 \times 2^5 + 0 \times 2^4 + 1 \times 2^3 + 0 \times 2^2 + 1 \times 2^1 + 0 \times 2^0 = 32 + 0 + 8 + 0 + 2 + 0 = 42$$

Exercise 24. (p. 16) After issuing `:set base=2`, here are Cryptol's responses:

```
0b1100
0b11000
0b1100
0b101100
0b110001
0b100000
0b1100100000
True
```

Exercise 25. (p. 16) `0` has the type `{a} [a]`. Incidentally, `0` is the only value that inhabits this type. The type `[2]` is precisely inhabited by the elements `0`, `1`, `2`, and `3`.

Exercise 26. (p. 16) The number 42 needs at least 6 bits to represent; hence the last expression fails. Note how type-inference helps, as users can give type annotations only when they need to be more specific. Unlike the C example, Cryptol will statically make sure that there will not be any overflow.

Exercise 27. (p. 17) Remember that Cryptol is big-endian and hence `12:[6]` is precisely `[False, False, True, True, False, False]`. Here are Cryptol's responses:

```

Cryptol> take' {3} 0xFF
7
Cryptol> take' {3} (12:[6])
1
Cryptol> drop' {3} (12:[6])
4
Cryptol> split' {3} (12:[6])
[0, 3, 0]
Cryptol> groupBy' {3} (12:[6])
[1, 4]

```

For instance, the expression `take' {3} (12:[6])` evaluates as follows:

```

take' {3} (12:[6])
= take (3, [False, False, True, True, False, False])
= [False, False, True]
= 1

```

Follow similar lines of reasoning to justify the results for the remaining expressions.

Exercise 28. (p. 17) Because of the leading zeros in `12:[12]`, they all produce different results:

```

Cryptol> take' {3} (12:[12])
0
Cryptol> drop' {3} (12:[12])
12
Cryptol> split' {3} (12:[12])
[0, 0, 12]
Cryptol> groupBy' {3} (12:[12])
[0, 0, 1, 4]

```

We will show the evaluation steps for `groupBy` here, and urge the reader to do the same for `splitBy`:

```

groupBy' {3} (12:[12])
= groupBy' {3} [False, False, False, False, False, False,
                False, False, True, True, False, False]
= [[False, False, False], [False, False, False]
   [False, False, True], [True, False, False]]
= [0, 0, 1, 4]

```

Exercise 29. (p. 17) Here are Cryptol's responses:

```

3
48

```

Section 2.8 Records: Named collections (p.18)

Exercise 30. (p. 18) Here are Cryptol's responses:

```

Cryptol> {xCoord = 12:[32], yCoord = 21:[32]}
{xCoord = 12, yCoord = 21}
Cryptol> {xCoord = 12:[32], yCoord = 21:[32]}.yCoord
21
Cryptol> {name = "Cryptol", address = "Galois"}
{name = "Cryptol", address = "Galois"}
Cryptol> {name = "Cryptol", address = "Galois"}.address
"Galois"
Cryptol> {name = "test", coords = {xCoord = 3:[32], yCoord = 5:[32]}}
{name = "test", coords = {xCoord = 3, yCoord = 5}}
Cryptol> {name = "test", coords = {xCoord = 3:[32], \
                                yCoord = 5:[32]}}.coords.yCoord
5
Cryptol> {x=True, y=False} == {y=False, x=True}
True

```

Section 2.9 The zero (p.19)

Exercise 31. (p. 19) Here are Cryptol's responses:

```

Cryptol> (zero : ([8] -> [3])) 5
0
Cryptol> (zero : Bit -> {xCoord : [12], yCoord : [5]}) True
{xCoord=0, yCoord=0}

```

The zero function returns 0, ignoring its argument.

Section 2.10 Arithmetic (p.20)

Exercise 32. (p. 20) Since 1 requires only 1-bit to represent, the result also has 1-bits. In other words, the arithmetic is done modulo $2^1 = 2$. Therefore, $1+1 = 0$.

Exercise 33. (p. 20) Now we have 8-bits to work with, so the result is 2. Since we have 8-bits to work with, overflow will not happen until we get a sum that is at least 256.

Exercise 34. (p. 20) Recall from Section 2.3 that there are no negative numbers in Cryptol. The values 3 and 5 can be represented in 3 bits, so Cryptol uses 3-bits to represent the result, so the arithmetic is done modulo $2^3 = 8$. Hence, the result is 6. In the second expression, we have 8-bits to work with, so the modulus is $2^8 = 256$; so the subtraction results in 254 (or 0xfe).

Exercise 35. (p. 20) The division/modulus by zero will give the expected error messages. In the last expression, the number 25 fits in 5 bits, so the modulus is $2^5 = 32$. The unary-minus yields 7, hence the result is 3. Note that \lg_2 is the *floor log base 2* function. The `width` function is the *ceiling log base 2* function.

Exercise 36. (p. 20) Here are Cryptol's answers:

```
(2, 0)
(2, 1)
(2, 2)
(3, 0)
```

The following equation holds regarding `/` and `%`:

$$x = (x/y) * y + (x\%y)$$

whenever $y \neq 0$.

Exercise 37. (p. 20) The bit-width in this case is 3 (to accommodate for the number 5), and hence arithmetic is done modulo $2^3 = 8$. Thus, `-2` evaluates to 6, leading to the result `min 5, (-2) == 5`. The parentheses are necessary because unary negation is handled in Cryptol's parser, not in its lexer, because whitespace is ignored. If this were not the case, reflect upon how you would differentiate the expressions `min 5 - 2` and `min 5 -2`.

Exercise 38. (p. 20) This time we are telling Cryptol to use precisely 8 bits, so `-2` is equivalent to 254. Therefore the result is 254.

Exercise 39. (p. 20) The idiomatic Cryptol way of summing two sequences is to use a comprehension:

```
[ i+j | i <- [1 .. 10]
      | j <- [10 .. 1]
]
```

However, you will notice that the following will work as well:

```
[1 .. 10] + [10 .. 1]
```

That is, Cryptol automatically lifts arithmetic operators to sequences, element-wise. However, it is often best to keep the explicit style of writing the comprehension, even though it is a bit longer, since that makes it absolutely clear what the intention is and how the new sequence is constructed, without depending implicitly upon Cryptol's automatic lifting.

Exercise 41. (p. 21) Here are Cryptol's responses:

```
[0]
[0, 0, 0, 0, 0, ...]
```

as opposed to `[0, 1, 0, 1, 0 ..]`, as one might expect⁶. This behavior follows from the specification that the width of the elements of the sequence are derived from the width of the elements in the seed, which in this case is 0.

Exercise 42. (p. 21) The expression `[1 .. 10]` is equivalent to `[1, (1+1) .. 10]`, and Cryptol knows that 10 requires at least 4-bits to represent and uses the minimum implied by all the available information. Hence we get: `[1, 2, 3, 4, 5, 6, 7, 8, 9, 10]`. You can use the `:t` command to see the type Cryptol infers for this expression explicitly:

⁶This is one of the subtle changes from Cryptol 1. The previous behavior can be achieved by dropping the first element from `[1 ..]`.

```
Cryptol> :t [1 .. 10]
{a} (a >= 4, fin a) => [10][a]
```

Cryptol tells us that the sequence has precisely 10 elements, and each element is at least 4-bits wide.

Section 2.11 Types (p.22)

Exercise 43. (p. 22) We have 12 elements, each of which is a sequence of 3 elements; so we have $12 * 3 = 36$ elements total. Each element is a 6-bit word; so the total number of bits is $36 * 6 = 216$.

Exercise 44. (p. 22) `[inf][inf][32]`. The size of such a value would be infinite!

Exercise 45. (p. 24) Here is the type of `groupBy`:

```
Cryptol> :t groupBy
groupBy : {each, parts, elem}
         (fin each) => [parts * each]elem
         -> [parts][each]elem
```

At every use case of `groupBy` we must instantiate the parameters `each`, `parts`, and `elem`; so that the resulting instantiation will match the use case. In the first example, we can simply take: `each = 3`, `parts = 3`, and `elem = [4]`. In the second, we can take `each=3`, `parts=4`, and `elem=[4]`. The third expression does not type check. Cryptol tells us:

```
Cryptol> groupBy' {3} [1..10] : [3][2][8]
Type mismatch:
  Expected type: 2
  Inferred type: 3
```

In this case, we are telling Cryptol that `each = 3`, `parts = 2`, and `elem = [8]` by providing the explicit type signature. Using this information, Cryptol must ensure that the instantiation will match the polymorphic type. To do so, Cryptol divides 10 (the size of the second argument) by 3 (the value of `each`) to obtain 3, and finds out that it does not match what we told it to use for `parts`, i.e., 2. It is not hard to see that there is no instantiation to make this work, since 10 is not divisible by 3.

The message we get for the last equation is truly interesting:

```
Cryptol> groupBy' {3} [1..10]
<polymorphic value>
Cryptol> :t groupBy' {3} [1..10]
{a, b} (a >= 4, fin a, 10 == 3 * b) => [b][3][a]
```

Cryptol is telling us that the result is a polymorphic value, for all values of `a` such that $3*a$ is 10. Type inference in the presence of arbitrary expressions is undecidable, and hence Cryptol tells us that this value will be instantiated to a concrete type as soon as we tell it what that `a` must be. Since there is no such `a`, we will never be able to use this value in a monomorphic context.

Exercise 46. (p. 25) Here is one way of writing this predicate, following the fact that $128 = 2 * 64$, $192 = 3 * 64$, and $256 = 4 * 64$:

```
{k} (2 <= k, k <= 4) => [k*64]
```

Here is another way, more direct but somewhat less satisfying:

```
{k} ((k - 128) * (k - 192) * (k - 256) == 0) => [k]
```

Note that Cryptol's type constraints do not include *or* predicates, hence we cannot just list the possibilities in a list.

Section 2.12 Defining functions (p.26)

Exercise 47. (p. 26) Here are some example uses of increment:

```
Cryptol> increment 3
4
Cryptol> increment 255
0
Cryptol> increment 912
[error] at <interactive>:1:1--1:14:
  Unsolved constraint:
    8 >= 10
    arising from use of expression demote at <interactive>:1:11--1:14
```

Note how type inference rejects application when applied to an argument of the wrong size: 912 is too big to fit into 8 bits.

Exercise 48. (p. 27) The signature indicates that `twoPlusXY` is a function that takes two 8-bit words as a tuple, and returns an 8-bit word.

Exercise 49. (p. 27) Here is the type Cryptol infers:

```
Cryptol> :t twoPlusXY
twoPlusXY : {a} (a >= 2, fin a) => ([a],[a]) -> [a]
```

That is, our function will actually work over arbitrary (finite) sized words, as long as they are at least 2-bits wide. The 2-bit requirement comes from the constant 2, which requires at least 2 bits to represent.

Exercise 50. (p. 27) Here is one way of defining this function:

```
minMax4 : {a} (Cmp a) => [4]a -> (a, a)
minMax4 [a, b, c, d] = (e, f)
  where e = min a (min b (min c d))
        f = max a (max b (max c d))
```

Note that ill-typed arguments will be caught at compile time! So, the second invocation with the 5 element sequence will fail to type-check. The `Cmp a` constraint arises from the types of `min` and `max` primitives:

```
min, max : {a} (Cmp a) => a -> a -> a
```


Exercise 51. (p. 27) Using `reverse` and `tail`, `butLast` is easy to define:

```
butLast : {n, t} (fin n) => [n+1]t -> [n]t
butLast xs = reverse (tail (reverse xs))
```

Here is another way to define `butLast`:

```
butLast' : {count, x} (fin count) => [count+1]x -> [count]x
butLast' xs = take' {count} xs
```

The type signature sets `count` to the desired width of the output, which is one shorter than the width of the input:

```
Cryptol> butLast []

[error] at <interactive>:1:1--1:11:
Unsolved constraint:
  0 >= 1
  arising from
  matching types
  at <interactive>:1:1--1:11
```

At first the error message might be confusing. What Cryptol is telling us that it deduced `count+1` must be 1, which makes `count` have value 0. But the `count+1` we gave it was 0, which is not greater than or equal to 1.

Finally, note that `butLast` requires a finite sequence as input, for obvious reasons, and hence the `fin n` constraint.

Exercise 52. (p. 28)

```
all f xs = [ f x | x <- xs ] == ~zero
```

Note how we apply `f` to each element in the sequence and check that the result consists of all `Trues`, by using a complemented zero. If we pass any an empty sequence, then we will always get `False`:

```
Cryptol> any eqTen [] where eqTen x = x == 10
False
```

This is intuitively the correct behavior as well. The predicate is satisfied by *none* of the elements in the sequence, but `any` requires at least one.

Exercise 53. (p. 28)

```
any f xs = [ f x | x <- xs ] != zero
```

This time all we need to make sure is that the result is not zero, i.e., at least one of elements yielded `True`. If we pass `all` an empty sequence, then we will always get `False`:

```
Cryptol> any eqTen [] where eqTen x = x == 10
False
```

Again, this is the correct response since there are no elements in an empty sequence that can satisfy the predicate.

Section 2.13 Recursion and recurrences (p.28)

Exercise 54. (p. 28)

```
isOdd, isEven : {n} (fin n, n >= 1) => [n] -> Bit
isOdd x = if x == 0 then False else isEven (x - 1)
isEven x = if x == 0 then True  else isOdd  (x - 1)
```

The extra predicate we need to add is $n \geq 1$. This constraint comes from the subtraction with 1, which requires at least 1 bit to represent.

Exercise 55. (p. 28) A number is even if its least significant bit is `False`, and odd otherwise. Hence, we can define these functions as:

```
isOdd', isEven' : {n} (fin n, n >= 1) => [n] -> Bit
isOdd' x = x ! zero
isEven' x = ~(x ! zero)
```

Note the use of `zero` which permits Cryptol to choose the width of the 0 constant appropriately.

Exercise 56. (p. 29) The type of `maxSeq` is:

```
maxSeq : {a, b} (fin a, fin b) => [a][b] -> [b]
```

It takes a sequence of words and returns a word of the same size. The suggested expressions produce 0, 10, and 10, respectively.

Exercise 57. (p. 29) We can simply drop the selection of the last element (`! 0`), and write `maxSeq'` as follows:

```
maxSeq' : {a, b} (fin a, fin b) => [a][b] -> [1 + a][b]
maxSeq' xs = ys
  where ys = [0] # [ max x y | x <- xs
                    | y <- ys
                ]
```

Exercise 58. (p. 30) Here is one answer. Note that in this solution the width of the answer is specified in terms of the width of the elements, so is likely to overflow. You can prevent the overflow by explicitly specifying the width of the output.

```
sumAll : {n, a} (fin n, fin a) => [n][a] -> [a]
sumAll xs = ys ! 0
  where ys = [0] # [ x+y | x <- xs
                    | y <- ys
                ]
```

In this code, the sequence `ys` contains the partial running sums. This is precisely the same pattern we have seen in Example 57. The output for the example calls are:

```
CrashCourse> sumAll []
0
CrashCourse> sumAll [1]
1
CrashCourse> sumAll [1, 2]
3
CrashCourse> sumAll [1, 2, 3]
6
CrashCourse> sumAll [1, 2, 3, 4]
10
CrashCourse> sumAll [1, 2, 3, 4, 5]
15
CrashCourse> sumAll [1 .. 100]
5050
```

If we do not explicitly tell Cryptol how wide the result is, then it will pick the width of the input elements, which will cause overflow and be subject to modular arithmetic as usual. Experiment with different signatures for `sumAll`, to avoid the overflow automatically, to get the answers:

```
0
1
3
6
10
15
5050
```

Exercise 59. (p. 31) Using a fold, it is easy to write `elem`:

```
elem (x, xs) = matches ! 0
  where matches = [False] # [ m || (x == e) | e <- xs
                        | m <- matches
                    ]
```

Note how we or (`||`) the previous result `m` with the current match, accumulating the result as we walk over the sequence. Starting with `False` ensures that if none of the matches succeed we will end up returning `False`. We have:

```
Cryptol> elem (2, [1..10])
True
Cryptol> elem (0, [1..10])
False
Cryptol> elem (10, [])
False
```

Exercise 60. (p. 31)

```

fibs : [inf][32]
fibs = [0, 1] # [ x+y | x <- fibs
                | y <- drop' {1} fibs
                ]

```

In this case we use the sequence `[0, 1]` as the seed, and both branches recursively refer to the defined value `fibs`. In the second branch, we drop the first element to skip over the first element of the sequence, effectively pairing the previous two elements at each step. The n th fibonacci number is obtained by the expression `fibs @ n`:

```

Cryptol> fibs @ 3
2
Cryptol> fibs @ 4
3
Cryptol> take' {10} fibs
[0, 1, 1, 2, 3, 5, 8, 13, 21, 34]

```

Note that `fibs` is an infinite stream of 32 bit numbers, so it will eventually be subject to wrap-around due to modular arithmetic.

Section 2.14 Stream equations (p.31)

Exercise 61. (p. 31)

```

xs input = [0x89, 0xAB, 0xCD, 0xEF] # new
where new = [ a ^ b ^ c | a <- as
                | b <- drop' {2} as
                | c <- input ]

```

Section 2.15 Type synonyms (p.32)

Exercise 62. (p. 33) A point is on the a^{th} axis if its non- a^{th} components are 0. Hence we have:

```

type Point3D a = {x : [a], y : [a], z : [a]}

onAnAxis : {a} (fin a) => Point3D a -> Bit
onAnAxis p = onX || onY || onZ
  where onX = (p.y == 0) && (p.z == 0)
        onY = (p.x == 0) && (p.z == 0)
        onZ = (p.x == 0) && (p.y == 0)

```

Exercise 63. (p. 33) This code:

```
cmpArith x y z = if x == y then z else z+z
```

yields the inferred type:

```
cmpArith : {a, b} (Arith b, Cmp a) => a -> a -> b -> b
```

Section 3.1 Caesar's cipher (p.38)

Exercise 1. (p. 38) Here is the alphabet and the corresponding shift-2 Caesar's alphabet:

```
Cryptol> ['A'..'Z']  
"ABCDEFGHIJKLMNOPQRSTUVWXYZ"  
Cryptol> ['A'..'Z'] <<< 2  
"CDEFGHIJKLMNOPQRSTUVWXYZAB"
```

We use a left rotate to get the characters lined up correctly, as illustrated above.

Exercise 2. (p. 38) Here are Cryptol's responses:

```
Cryptol> caesar (0, "ATTACKATDAWN")  
"ATTACKATDAWN"  
Cryptol> caesar (3, "ATTACKATDAWN")  
"DWWDFNDWGDZQ"  
Cryptol> caesar (12, "ATTACKATDAWN")  
"MFFMOWMFPMIZ"  
Cryptol> caesar (52, "ATTACKATDAWN")  
"ATTACKATDAWN"
```

If the shift is a multiple of 26 (as in 0 and 52 above), the letters will cycle back to their original values, so encryption will leave the message unchanged. Users of the Caesar's cipher should be careful about picking the shift amount!

Exercise 3. (p. 38) The code is almost identical, except we need to use a right rotate:

```
dCaesar : {n} ([8], String n) -> String n  
dCaesar (s, msg) = [ shift x | x <- msg ]  
  where map      = ['A' .. 'Z'] >>> s  
        shift c = map @ (c - 'A')
```

We have:

```
Cryptol> caesar (12, "ATTACKATDAWN")  
"MFFMOWMFPMIZ"  
Cryptol> dCaesar (12, "MFFMOWMFPMIZ")  
"ATTACKATDAWN"
```

Exercise 4. (p. 38) For the Caesar’s cipher, the only good shifts are 1 through 25, since shifting by 0 would return the plaintext unchanged, and any shift amount d that is larger than 26 and over is essentially the same as shifting by $d \pmod{26}$ due to wrap around. Therefore, all it takes to break the Caesar cipher is to try the sizes 1 through 25, and see if we have a valid message. We can automate this in Cryptol by returning all possible plaintexts using these shift amounts:

```
attackCaesar : {n} (String n) -> [25](String n)
attackCaesar msg = [ dCaesar(i, msg) | i <- [1 .. 25] ]
```

If we apply this function to JHLZHYJPWOLYPZDLHR, we get:

```
Cryptol> :set ascii=on
Cryptol> attackCaesar "JHLZHYJPWOLYPZDLHR",
["IGKYGXIOVNKXOYCKGQ", "HFJXFVHNUMJWNXBJFF", "GEIWEVGMTLIVMWAIEO"
"FDHVDUFLSKHULVZHDN", "ECGUCTEKRJGTKUYGCM", "DBFTBSDJQIFSJTXFBL"
"CAESARCIPHERISWEAK", "BZDRZQBHOGDQHRVDZJ", "AYCOYPAGNFPCPGQUCYI"
"ZXPXOZFMEOFPBTBXH", "YWAOWNYELDANEOSAWG", "XVZNVMDKDCZMDNRZVF"
"WUYMULWCJBYLCMOYUE", "VTXLTKVBIAXKBLPXTD", "USWKSJUAHZWJAKOWSC"
"TRVJRITZGYVIZJNVRB", "SQUIQHSYFXUHYIMUQA", "RPTHPRXEWGTXHLTPZ"
"QOSGOFQWDVSWFGKSOY", "PNRFNEPVCUREVVFJRNX", "OMQEMDOUBTQDUEIQMW"
"NLPDLCNTASPCTDHPV", "MKOCKBMSZROBSCGOKU", "LJNBALRYQNARBFNJT"
"KIMAIZKQXPMZQAEMIS"]
```

If you skim through the potential ciphertexts, you will see that the 7th entry is probably the one we are looking for. Hence the key must be 7. Indeed, the message is CAESARCIPHERISWEAK.

Exercise 5. (p. 38) No. Using two shifts d_1 and d_2 is essentially the same as using just one shift with the amount $d_1 + d_2$. Our attack function would work just fine on this schema as well. In fact, we wouldn’t even have to know how many rounds of encryption was applied. Multiple rounds is just as weak as a single round when it comes to breaking the Caesar’s cipher.

Exercise 6. (p. 38) In this case we will fail to find a mapping:

```
Cryptol> caesar (3, "12")
... index of 240 is out of bounds
(valid range is 0 thru 25).
```

What happened here is that Cryptol computed the offset ‘1’ - ‘A’ to obtain the 8-bit index 240 (remember, modular arithmetic!), but our alphabet only has 26 entries, causing the out-of-bounds error. We can simply remedy this problem by allowing our alphabet to contain all 8-bit numbers:

```
caesar' : {n} ([8], String n) -> String n
caesar' (s, msg) = [ shift x | x <- msg ]
  where map      = [0 .. 255] <<< s
        shift c = map @ c
```

Note that we no longer have to subtract ‘A’, since we are allowing a much wider range for our plaintext and ciphertext. (Another way to put this is that we are subtracting the value of the first element in the alphabet, which happens to be 0 in this case! Consequently, the number of “good” shifts increase from 25 to 255.) The change in `dCaesar'` is analogous:

```

dCaesar' : {n} ([8], String n) -> String n
dCaesar' (s, msg) = [ shift x | x <- msg ]
  where map      = [0 .. 255] >>> s
        shift c = map @ c

```

Section 3.2 Vigenère cipher (p.38)

Exercise 7. (p. 39) Here is one way to define `cycle`, using a recursive definition:

```

cycle xs = xss
  where xss = xs # xss

```

We have:

```

Cryptol> cycle [1 .. 3]
[1, 2, 3, 1, 2, ...]

```

If we do not have the `n >= 1` predicate, then we can pass `cycle` the empty sequence, which would cause an infinite loop emitting nothing. The predicate `n >= 1` makes sure the input is non-empty, guaranteeing that `cycle` can produce the infinite sequence.

Exercise 8. (p. 39)

```

vigenere (key, pt) = join [ caesar (k - 'A', [c])
                          | c <- pt
                          | k <- cycle key
                          ]

```

Note the shift is determined by the distance from the letter 'A' for each character. Here is the cipher in action:

```

Cryptol> vigenere ("CRYPTOL", "ATTACKATDAWN")
"CKRPVYLVUULG"

```

Exercise 9. (p. 39) Following the lead of the encryption, we can rely on `dCaesar`:

```

dVigenere : {n, m} (fin n, n >= 1) =>
  (String n, String m) -> String m
dVigenere (key, pt) = join [ dCaesar (k - 'A', [c])
                          | c <- pt
                          | k <- cycle key
                          ]

```

The secret code is:

```

Cryptol> dVigenere ("CRYPTOL", "XZETGSCGTYCMGEQGAGRDEQC")
"VIGENERECANTSTOPCRYPTOL"

```

Exercise 10. (p. 39) All it takes is to decrypt using the plaintext as the key and message as the cipherkey. Here is this process in action. Recall from the previous exercise that encrypting `ATTACKATDAWN` by the key `CRYPTOL` yields `CKRPVYLVUYLG`. Now, if an attacker knows that `ATTACKATDAWN` and `CKRPVYLVUYLG` form a pair, he/she can find the key simply by:

```
Cryptol> dvigenere ("ATTACKATDAWN", "CKRPVYLVUYLG")
"CRYPTOLCRYPT"
```

Note that this process will not always tell us what the key is precisely. It will only be the key repeated for the given message size. For sufficiently large messages, or when the key does not repeat any characters, however, it would be really easy for an attacker to glean the actual key from this information.

This trick works since the act of using the plaintext as the key and the ciphertext as the message essentially reverses the shifting process, revealing the shift amounts for each pair of characters. The same attack would essentially work for the Caesar's cipher as well, where we would only need one character to crack it.

Section 3.3 The atbash (p.39)

Exercise 11. (p. 39) Using the reverse index operator, coding atbash is trivial:

```
atbash : {n} String n -> String n
atbash pt = [ alph ! (c - 'A') | c <- pt ]
  where alph = ['A' .. 'Z']
```

We have:

```
Cryptol> atbash "ATTACKATDAWN"
"ZGGZXPZGWZDM"
```

Exercise 12. (p. 39) Notice that decryption for atbash is precisely the same as encryption, the process is entirely the same. So, we do not have to write any code at all, we can simply define:

```
dAtbash : {n} String n -> String n
dAtbash = atbash
```

We have:

```
Cryptol> dAtbash "ZGYZHSRHHVOUVXIBKGRMT"
"ATBASHISSELFDECRYPTING"
```

Section 3.4 Substitution ciphers (p.40)

Exercise 13. (p. 40)

```
subst (key, pt) = [ key @ (p - 'A') | p <- pt ]
```


We have:

```
Cryptol> subst(substKey, "SUBSTITUTIONSSAVETHEDAY")
"NLJNUXULUXAINNFSOUROWFC"
```

Exercise 14. (p. 40)

```
invSubst (key, c) = candidates ! 0
  where candidates = [0] # [ if c == k then a else p
                            | k <- key
                            | a <- ['A' .. 'Z']
                            | p <- candidates
                            ]
```

The comprehension defining `candidates` uses a fold (see page 29). The first branch (`k <- key`) walks through all the key elements, the second branch walks through the ordinary alphabet (`a <- ['A' .. 'Z']`), and the final branch walks through the candidate match so far. At the end of the fold, we simply return the final element of `candidates`. Note that we start with 0 as the first element, so that if no match is found we get a 0 back.

Exercise 15. (p. 41)

```
dSubst: {n} (String 26, String n) -> String n
dSubst (key, ct) = [ invSubst (key, c) | c <- ct ]
```

We have:

```
Cryptol> dSubst (substKey, "FUUFHKFUWFGI")
"ATTACKATDAWN"
```

Exercise 16. (p. 41) No, with this key we cannot decrypt properly:

```
Cryptol> subst ("AAAABBBBCCCCDDDEEEFFFFGG", "HELLOWORLD")
"BBCCDFDECA"
Cryptol> dSubst ("AAAABBBBCCCCDDDEEEFFFFGG", "BBCCDFDECA")
"HLLLPXPTLD"
```

This is because the given key maps multiple plaintext letters to the same ciphertext letter. (For instance, it maps all of A, B, C, and D to the letter A.) For substitution ciphers to work the key should not repeat the elements, providing a 1-to-1 mapping. This property clearly holds for `substKey`. Note that there is no shortage of keys, since for 26 letters we have $26!$ possible ways to choose keys, which gives us over 4-billion different choices.

Section 3.5 The scytale (p.41)

Exercise 17. (p. 42) If you do not provide a signature for `msg'`, you will get the following type-error message from Cryptol:

```

Failed to validate user-specified signature.
In the definition of 'scytale', at classic.cry:40:1--40:8:
  for any type row, diameter
    fin row
    fin diameter
=>
fin ?b
  arising from use of expression split at classic.cry:42:17--42:22
fin ?d
  arising from use of expression join at classic.cry:40:15--40:19
row * diameter == ?a * ?b
  arising from matching types at classic.cry:1:1--1:1

```

Essentially, Cryptol is complaining that it was asked to do a `split` and it figured that the constraint $diameter * row = a * b$ must hold, but that is not sufficient to determine what `a` and `b` should really be. (There could be multiple ways to assign `a` and `b` to satisfy that requirement, for instance `a=4, b=row`; or `a=2` and `b=2*row`, resulting in differing behavior.) This is why it is unable to “validate the user-specified signature”. By putting the explicit signature for `msg'`, we are giving Cryptol more information to resolve the ambiguity. Notice that since the code for `scytale` and `dscytale` are precisely the same except for the type on `msg'`. This is a clear indication that the type signature plays an essential role here.

Exercise 18. (p. 42) Even if we do not know the diameter, we do know that it is a divisor of the length of the message. For any given message size, we can compute the number of divisors of the size and try decryption until we find a meaningful plaintext. Of course, the number of potential divisors will be large for large messages, but the practicality of `scytale` stems from the choice of relatively small diameters, hence the search would not take too long. (With large diameters, the ancient Greeks would have to carry around very thick rods, which would not be very practical in a battle scenario!)

Section 4.1 The plugboard (p.43)

Exercise 1. (p. 44) We can simply ask Cryptol what the implied mappings are:

```

Cryptol> [ plugboard @ (c - 'A') | c <- "ACQUTUWO" ]
"HGXVYML"

```

Why do we subtract the 'A' when indexing?

Section 4.2 Scrambler rotors (p.44)

Exercise 2. (p. 44) Recall that `rotor1` was defined as:

```

rotor1 = mkRotor ("RJICAWVQZODLUPYFEHXSMTKNGB", "IO")

```

Here is a listing of the new mappings and the characters we will get at the output for each successive `c`:

starting map	output	notch engaged?
RJICAWVQZODLUPYFEHXSMTKNGB	I	no
JICAWVQZODLUPYFEHXSMTKNGBR	C	no
ICAWVQZODLUPYFEHXSMTKNGBRJ	A	yes
CAWVQZODLUPYFEHXSMTKNGBRJI	W	no
AWVQZODLUPYFEHXSMTKNGBRJIC	V	no

Note how we get different letters as output, even though we are providing the same input (all c's.) This is the essence of the Enigma: the same input will not cause the same output necessarily, making it a polyalphabetic substitution cipher.

Section 4.3 Connecting the rotors: notches in action (p.45)

Exercise 3. (p. 45) We can define the following value to simulate the operation of always telling `scramble` to rotate the rotor and providing it with the input `c`.

```
rotor1CCCC = [(c1, n1), (c2, n2), (c3, n3), (c4, n4), (c5, n5)]
  where (n1, c1, r1) = scramble (True, 'C', rotor1)
        (n2, c2, r2) = scramble (True, 'C', r1)
        (n3, c3, r3) = scramble (True, 'C', r2)
        (n4, c4, r4) = scramble (True, 'C', r3)
        (n5, c5, r5) = scramble (True, 'C', r4)
```

Note how we chained the output rotor values in the calls, through the values `r1-r2-r3` and `r4`. We have:

```
Cryptol> rotor1CCCC
[(I, False), (C, False), (A, True), (W, False), (V, False)]
```

Note that we get precisely the same results from `Cryptol` as we predicted in the previous exercise.

Exercise 4. (p. 47) Not unless we receive an empty sequence of rotors, i.e., a call of the form: `joinRotors ([], c)` for some character `c`. In this case, it does make sense to return `c` directly, which is what `initRotor` will do. Note that unless we do receive an empty sequence of rotors, the value of `initRotor` will not be used when computing the `joinRotors` function.

Exercise 5. (p. 47) The crucial part is the value of `ncrs`. Let us write it out by substituting the values of rotors and `inputChar`:

```
ncrs = [(True, 'F', initRotor)]
      # [ scramble (notch, char, r)
        | r <- [rotor1, rotor2, rotor3]
        | (notch, char, rotor') <- ncrs
        ]
```

Clearly, the first element of `ncrs` will be:

```
(True, 'F', initRotor)
```

Therefore, the second element will be the result of the call:

```
scramble (True, 'F', rotor1)
```

Recall that `rotor1` was defined as:

```
rotor1 = mkRotor ("RJICAWVQZODLUPYFEHXSMTKNGB", "IO")
```

What letter does `rotor1` map `F` to? Since `F` is the 5th character (counting from 0), `rotor1` maps it to the 5th element of its permutation, i.e., `w`, remembering to count from 0! The topmost element in `rotor1` is `R`, which is not its notch-list, hence it will *not* tell the next rotor to rotate. But it will rotate itself, since it received the `True` signal. Thus, the second element of `ncrs` will be:

```
(False, 'W', ...)
```

where we used `...` to denote the one left-rotation of `rotor1`. (Note that we do not need to know the precise arrangement of `rotor1` now for the purposes of this exercise.) Now we move to `rotor2`, we have to compute the result of the call:

```
scramble (False, 'W', rotor2)
```

Recall that `rotor2` was defined as:

```
rotor2 = mkRotor ("DWYOLETKNVQPHURZJMSFIGXCBA", "B")
```

So, it maps `w` to `x`. (The fourth letter from the end.) It will not rotate itself, and it will not tell `rotor3` to rotate itself either since the topmost element is `D` in its current configuration, and `D` which is not in the notch-list `"B"`. Thus, the final `scramble` call will be:

```
scramble (False, 'X', rotor3)
```

where

```
rotor3 = mkRotor ("FGKMAJWUOVNRYIZETDPSHBLCOX", "CK")
```

It is easy to see that `rotor3` will map `x` to `c`. Thus the final value coming out of this expression must be `c`. Indeed, we have:

```
Cryptol> project(2, 2, joinRotors ([rotor1 rotor2 rotor3], 'F'))
c
```

Of course, Cryptol also keeps track of the new rotor positions as well, which we have glossed over in this discussion.

Section 4.4 The reflector (p.47)

Exercise 6. (p. 47)

```
all : {n, a} (fin n) => (a -> Bit) -> [n]a -> Bit
all f xs = [ f x | x <- xs ] == ~zero
```

```
checkReflector refl = all check ['A' .. 'Z']
  where check c = (c != m) && (c == c')
        where m = refl @ (c - 'A')
              c' = refl @ (m - 'A')
```

For each character in the alphabet, we first figure out what it maps to using the reflector, named `m` above. We also find out what `m` gets mapped to, named `c'` above. To be a valid reflector it must hold that `c` is not `m` (no character maps to itself), and `c` must be `c'`. We have:

```
Cryptol> checkReflector reflector
True
```

Note how we used `all` to make sure `check` holds for all the elements of the alphabet.

Section 4.5 Putting the pieces together (p.47)

Exercise 7. (p. 48) We can define the following helper function, using the function `all` you have defined in Exercise 2.9-52:

```
checkPermutation : Permutation -> Bit
checkPermutation perm = all check ['A' .. 'Z']
  where check c = (c == substBwd(perm, substFwd (perm, c)))
                && (c == substFwd(perm, substBwd (perm, c)))
```

Note that we have to check both ways (first `substFwd` then `substBwd`, and also the other way around) in case the substitution is badly formed, for instance if it is mapping the same character twice. We have:

```
Cryptol> checkPermutation [ c | (c, _) <- rotor1 ]
True
```

For a bad permutation we would get `False`:

```
Cryptol> checkPermutation (['A' .. 'Y'] # ['A'])
False
```

Exercise 8. (p. 48) Since the reflector is symmetric, substituting backwards or forwards does not matter. We can verify this with the following helper function:

```
all : {a, b} (fin b) => (a -> Bit) -> [b]a -> Bit
all fn xs = folds ! 0 where
  folds = [True] # [ fn x && p | x <- xs
                  | p <- folds]
checkReflectorFwdBwd : Reflector -> Bit
checkReflectorFwdBwd refl = all check ['A' .. 'Z']
  where check c = substFwd (refl, c) == substBwd (refl, c)
```

We have:

```
Cryptol> checkReflectorFwdBwd reflector
True
```

Section 4.7 Encryption and decryption (p.49)

Exercise 10. (p. 50) Enigma will start repeating once the rotors go back to their original position. With n rotors, this will take 26^n characters. In the case of the traditional 3-rotor Enigma this amounts to $26^3 = 17576$ characters. Note that we are assuming an ideal Enigma here with no double-stepping [2].

Exercise 11. (p. 50) Since the period for the 3-rotor Enigma is 17576 (see the previous exercise), we need to make sure two instances of CRYPTOL are 17576 characters apart. Since CRYPTOL has 7 characters, we should have 17569 X's. The following Cryptol definition would return the relevant pieces:

```
enigmaCryptol = (take' {7} ct, drop' {17576} ct)
  where   str = "CRYPTOL" # [ 'X' | _ <- [1 .. 17569] ]
         # "CRYPTOL"
         ct = dEnigma(modelEnigma, str)
```

We have:

```
Cryptol> enigmaCryptol
("KGSHPMK", "KGSHPMK")
```

As predicted, both instances of CRYPTOL get encrypted as KGSHPMK.

Section 5.1 Writing properties (p.51)

Exercise 1. (p. 51) Cryptol will print the property location, name, and the type. The command `:i` stands for `info`. It provides data about the properties, type-synonyms, etc. available at the top-level of your program.

Exercise 2. (p. 52)

```
property revRev xs = reverse (reverse xs) == xs
```

Exercise 3. (p. 52)

```
property appAssoc (xs, ys, zs) = xs # (ys # zs) == (xs # ys) # zs
```

Exercise 4. (p. 52)

```
property revApp (xs, ys) = reverse (xs # ys)
  == reverse ys # reverse xs
```

Exercise 5. (p. 52)

```
property lshMul (n, k) = n << k == n * 2^^k
```

Exercise 6. (p. 54) There are many such types, all sharing the property that they do not take any space to represent. Here are a couple examples:

```
Cryptol> flipNeverIdentity (zero : ([0], [0]))
False
Cryptol> flipNeverIdentity (zero : [0][8])
False
```

Exercise 7. (p. 54)

```
Cryptol> :t widthPoly
widthPoly : {a, b} (fin a) => [a]b -> Bit
```

It is easy to see that `widthPoly` holds at the instances:

```
{b} [15]b -> Bit
```

and

```
{b} [531]b -> Bit
```

but at no other. Based on this, we can write `evenWidth` as follows:

```
property evenWidth x = (width x) ! 0 == False
```

remembering that the 0th bit of an even number is always `False`. We have:

```
Cryptol> evenWidth (0:[1])
False
Cryptol> evenWidth (0:[2])
True
Cryptol> evenWidth (0:[3])
False
Cryptol> evenWidth (0:[4])
True
Cryptol> evenWidth (0:[5])
False
```

Section 5.2 Establishing correctness (p.54)

Exercise 8. (p. 56) If we try to prove `revRev` directly, we will get an error from Cryptol:

```
Cryptol> :prove revRev
Not a valid predicate type:
{a, b} (fin a, Cmp b) => [a]b -> Bit
```

Cryptol is telling us that the property has a polymorphic type, and hence cannot be proven. We can easily prove instances of it, by either creating new properties with fixed type signatures, or by monomorphising it via type annotations. Several examples are given below:

```
Cryptol> :prove revRev : [10][8] -> Bit
Q.E.D.
Cryptol> :prove revRev : [100][32] -> Bit
Q.E.D.
Cryptol> :prove revRev : [0][4] -> Bit
Q.E.D.
```

Exercise 12. (p. 56) We have:

```
Cryptol> :prove widthPoly : [15] -> Bit
Q.E.D.
Cryptol> :prove widthPoly : [531] -> Bit
Q.E.D.
Cryptol> :prove widthPoly : [8] -> Bit
widthPoly:[8] -> Bit 0 = False
Cryptol> :prove widthPoly : [32] -> Bit
widthPoly:[32] -> Bit 0 = False
```

Exercise 13. (p. 57)

```
property divModMul (x,y) = if y == 0
                          then True // precondition fails => True
                          else x == (x / y) * y + x % y
```

We have:

```
Cryptol> :prove divModMul : ([4], [4]) -> Bit
Q.E.D.
Cryptol> :prove divModMul : ([8], [8]) -> Bit
Q.E.D.
```

Exercise 14. (p. 57) Using `all` and `elem`, it is easy to express `validMessage`:

```
validMessage = all (\c -> elem (c, ['A' .. 'Z']))
```

Note the use of a λ -expression to pass the function to `all`. Of course, we could have defined a separate function for it in a `where`-clause, but the function is short enough to directly write it inline.

Exercise 15. (p. 57) A naive attempt would be to write:

```
property caesarCorrectBOGUS (d,msg) =
  dCaesar(d, caesar(d, msg)) == msg
```


However, this property is not correct for all `msg`'s, since Caesar's cipher only works for messages containing the letters 'A' ... 'Z', not arbitrary 8-bit values as the above property suggests. We can see this easily by providing a bad input:

```
Cryptol> caesar (3, "1")
invalid sequence index: 240
```

(240 is the difference between the ASCII value of '1', 49, and the letter 'A', 65, interpreted as an 8-bit offset.) We should use the `validMessage` function of the previous exercise to write a conditional property instead:

```
property caesarCorrect (d,msg) = if validMessage msg
                                then dCaesar(d, caesar(d, msg)) == msg
                                else True
```

We have:

```
Cryptol> :prove caesarCorrect : ([8], String(10)) -> Bit
Q.E.D.
```

Exercise 16. (p. 57)

```
property modelEnigmaCorrect pt =
  if validMessage pt
  then dEnigma (modelEnigma, enigma (modelEnigma, pt)) == pt
  else True
```

We have:

```
Cryptol> :prove modelEnigmaCorrect : String(10) -> Bit
Q.E.D.
```

Section 5.3 Automated random testing (p.57)

Exercise 17. (p. 58) Here is the interaction with Cryptol (when you actually run this, you will see the test cases counting up as they are performed):

```
Cryptol> :check (caesarCorrect : ([8], String 10) -> Bit)
Using random testing.
passed 100 tests.
Coverage: 0.00% (100 of 288 values)
Cryptol> :set tests=1000
Cryptol> :check (caesarCorrect : ([8], String 10) -> Bit)
Using random testing.
passed 1000 tests.
Coverage: 0.00% (1000 of 288 values)
```

In each case, Cryptol tells us that it checked a minuscule portion of all possible test cases: A good reminder of what `:check` is really doing. The number of test cases is: $2^{8+8 \times 10} = 2^{88}$. We have 8-bits for the `d` value, and $10 * 8$ bits total for the 10 characters in `msg`, giving us a total of 88 bits. Since the input is 88 bits wide, we have 2^{88} potential test cases. Note how the number of test cases increase exponentially with the size of the message.

Exercise 18. (p. 58)

```
Cryptol> :check True
Using exhaustive testing.
passed 1 tests.
QED
Cryptol> :check False
Using exhaustive testing.
FAILED for the following inputs:
Cryptol> :check \x -> x == (x:[8])
Using exhaustive testing.
passed 256 tests.
QED
```

Note that when Cryptol is able to exhaust all possible inputs, it returns QED, since the property is effectively proven.

Exercise 19. (p. 58)

```
property easyBug x = x != (76123:[64])
```

The `:prove` command will find the counterexample almost instantaneously, while `:check` will have a hard time!

Section 5.4 Checking satisfiability (p.58)

Exercise 20. (p. 59)

```
modFermat (a, b, c, n) = (a > 0) && (b > 0) && (c > 0) && (n > 2)
&& (a^n + b^n == c^n)
```

The `fin s` predicate comes from the fact that we are doing arithmetic and comparisons. The predicate `s >= 2` comes from the fact that we are comparing `n` to 2, which needs at least 2 bits to represent.

Exercise 21. (p. 59) We can try different instantiations as follows:

```
Cryptol> :sat modFermat : Quad(2) -> Bit
modFermat : Quad(2) -> Bit (1, 2, 1, 3) = True
Cryptol> :sat modFermat : Quad(3) -> Bit
modFermat : Quad(3) -> Bit (4, 4, 4, 4) = True
Cryptol> :sat modFermat : Quad(4) -> Bit
modFermat : Quad(4) -> Bit (4, 4, 4, 8) = True
```



```

Cryptol> pdiv <| x^5 + x^3 + 1 |> <| x^3 + x^2 + 1 |> == \
  <| x^2 + x |>
True
Cryptol> pmod <| x^5 + x^3 + 1 |> <| x^3 + x^2 + 1 |> == \
  <| x^2 + x + 1 |>
True

```

Another way to check your result would be to multiply the quotient by the divisor and add the remainder, and check that it gives us precisely the polynomial we started with:

```

Cryptol> pmult <| x^2 + x |> <| x^3 + x^2 + 1 |> \
  ^ <| x^2 + x + 1 |>
0x29
Cryptol> <| x^5 + x^3 + 1 |>
0x29

```

Exercise 5. (p. 64)

```

property gf28MultUnit      x =      gf28Mult(x, 1) == x
property gf28MultCommutative x y =    gf28Mult(x, y) == gf28Mult(y, x)
property gf28MultAssociative x y z =  gf28Mult(x, gf28Mult(y, z))
                                == gf28Mult(gf28Mult(x, y), z)

```

It turns out that proving the unit and commutativity are fairly trivial:

```

Cryptol> :prove gf28MultUnit
Q.E.D.
Cryptol> :prove gf28MultCommutative
Q.E.D.

```

But `gf28MultAssociative` takes much longer! We show the results of `:check` below:

```

Cryptol> :check gf28MultAssociative
Checking case 1000 of 1000 (100.00%)
1000 tests passed OK

```

Note that the coverage is pretty small (on the order of 0.006%) in this case. Proving associativity of multiplication algorithms using SAT/SMT based technologies is a notoriously hard task [10, Section 6.3.1]. If you have the time, you can let Cryptol run long enough to complete the `:prove gf28MultAssociative` command, however.

Section 6.3 The `subBytes` transformation (p.64)

Exercise 6. (p. 64)

```

gf28Pow (n, k) = pow k
  where sq x = gf28Mult (x, x)
        odd x = x ! 0
        pow i = if i == 0 then 1

```

```

else if odd i
  then gf28Mult (n, sq (pow (i >> 1)))
  else sq (pow (i >> 1))

```

Here is a version that follows the stream-recursion pattern:

```

gf28Pow' : (GF28, [8]) -> GF28
gf28Pow' (n, k) = pows ! 0
  where pows = [1] # [ if bit then gf28Mult (n, sq x)
                       else sq x
                     | x <- pows
                     | bit <- k
                     ]
    sq x = gf28Mult (x, x)

```

Exercise 7. (p. 64)

```
gf28Inverse x = gf28Pow (x, 254)
```

We do not have to do anything special about 0, since our `gf28Inverse` function yields 0 in that case:

```

Cryptol> gf28Inverse 0
0x00

```

Exercise 8. (p. 65) Since 0 does not have a multiplicative inverse, we have to write a conditional property:

```

property gf28InverseCorrect x =
  if x == 0 then x' == 0 else gf28Mult(x, x') == 1
  where x' = gf28Inverse x

```

We have:

```

Cryptol> :prove gf28InverseCorrect
Q.E.D.

```

Exercise 9. (p. 65)

```

xformByte b = gf28Add [b, (b >>> 4), (b >>> 5),
                      (b >>> 6), (b >>> 7), c]
  where c = 0x63

```

Exercise 10. (p. 66)

```
property SubByteCorrect x = SubByte x == SubByte' x
```

We have:

```

Cryptol> :prove SubByteCorrect
Q.E.D.

```

Section 6.4 The `shiftRows` transformation (p.66)

Exercise 11. (p. 67) Consider what happens after 4 calls to `shiftRows`. The first row will stay the same, since it never moves. The second row moves one position each time, and hence it will move 4 positions at the end, restoring it back to its original starting configuration. Similarly, row 3 will rotate $2 \times 4 = 8$ times, again restoring it. Finally, row 4 rotates 3 times each for a total of $3 \times 4 = 12$ times, cycling back to its starting position. Hence, every 4th rotation will restore the entire state back. We can verify this in Cryptol by the following property:

```
shiftRow4RestoresBack : State -> Bit
property shiftRow4RestoresBack state = states @ 4 == state
  where states = [state] # [ ShiftRows s | s <- states ]
```

We have:

```
Cryptol> :prove shiftRow4RestoresBack
Q.E.D.
```

Of course, any multiple of 4 would have the same effect.

Section 6.5 The `mixColumns` transformation (p.67)

Exercise 12. (p. 67)

```
gf28DotProduct (xs, ys) = gf28Add [ gf28Mult (x, y) | x <- xs
                                   | y <- ys
                                   ]
```

Exercise 13. (p. 67)

```
property DPComm a b =   gf28DotProduct (a, b) == gf28DotProduct (b, a)
property DPDist a b c = gf28DotProduct (a, vectorAdd(b, c)) ==
  gf28Add [ab, ac]
  where ab = gf28DotProduct (a, b)
        ac = gf28DotProduct (a, c)
        vectorAdd (xs, ys) = [ gf28Add [x, y] | x <- xs
                               | y <- ys
                               ]
```

We have:

```
AES> :prove DPComm : [10]GF28 -> [10]GF28 -> Bit
Q.E.D.
```

```
AES> :check DPDist : [10]GF28 -> [10]GF28 -> [10]GF28 -> Bit
Using random testing.
passed 1000 tests.
```

You might be surprised that the total number of cases for this property is $2^{3 \cdot 10 \cdot 8} = 2^{240}$ —a truly ginormous number!

Exercise 14. (p. 67)

```
gf28VectorMult (v, ms) = [ gf28DotProduct(v, m) | m <- ms ]
```

Exercise 15. (p. 67) We simply need to call `gfVectorMult` of the previous exercise on every row of the first matrix, after transposing the second matrix to make sure columns are properly aligned. We have:

```
gf28MatrixMult (xss, yss) = [ gf28VectorMult(xs, yss')
                             | xs <- xss ]
  where yss' = transpose yss
```

Section 6.6 Key expansion (p.68)

Exercise 16. (p. 68) Finding out the elements is easy:

```
Cryptol> [ (Rcon i)@0 | i <- [1 .. 10] ]
[1, 2, 4, 8, 16, 32, 64, 128, 27, 54]
```

Note that we only capture the first element in each `Rcon` value, since we know that the rest are 0. We can now use this table to define `Rcon'` as follows:

```
Rcon' i = [(rcons @ (i-1)), 0, 0, 0]
  where rcons : [10]GF28
         rcons = [1, 2, 4, 8, 16, 32, 64, 128, 27, 54]
```

Note that we subtract 1 before indexing into the `rcons` sequence to get the indexing right.

Exercise 17. (p. 68) We need to write a conditional property (Section 5.2.4). Below, we use the function `elem` you have defined in Exercise 2.13-59:

```
property RconCorrect x = if elem(x, [1..10])
                          then Rcon x == Rcon' x
                          else True
```

We have:

```
Cryptol> :prove RconCorrect
Q.E.D.
```

Section 6.8 AES encryption (p.71)

Exercise 18. (p. 72)

```
property msgToStateToMsg msg = stateToMsg(msgToState(msg)) == msg
property stToMsgToSt s = msgToState(stateToMsg s) == s
```

We have:

```
Cryptol> :prove msgToStateToMsg
Q.E.D.
Cryptol> :prove stToMsgToSt
Q.E.D.
```

Section 6.9 Decryption (p.73)

Exercise 20. (p. 73)

```
property xformByteInverse x = xformByte' (xformByte x) == x
```

We have:

```
Cryptol> :prove xformByteInverse
Q.E.D.
```

Exercise 21. (p. 74)

```
property sbxInverse s = InvSubBytes (SubBytes s) == s
```

We have:

```
Cryptol> :prove xformByteInverse
Q.E.D.
```

Exercise 23. (p. 74)

```
property shiftRowsInverse s = InvShiftRows (ShiftRows s) == s
```

We have:

```
Cryptol> :prove shiftRowsInverse
Q.E.D.
```

Exercise 24. (p. 74)


```
property mixColumnsInverse s = InvMixColumns (MixColumns s) == s
```

Unlike others, this property is harder to prove automatically and will take much longer. Below we show the `:check` results instead:

```
Cryptol> :check mixColumnsInverse  
Checking case 1000 of 1000 (100.00%)  
1000 tests passed OK
```


Appendix B

Cryptol primitive functions

Bitwise operations

```
&&, ||, ^ : {a} a -> a -> a
~          : {a} a -> a
```

Comparisons

```
==, !=      : {a} (Cmp a) => a -> a -> Bit
<, >, <=, >= : {a} (Cmp a) => [a] -> [a] -> Bit
```

Arithmetic

```
+, -, *, /, %, ** : {a} (Arith a) => a -> a -> a
lg2               : {a, b} (Arith a) => a -> a
```

Polynomial arithmetic

```
pdiv : {a, b} (fin a, fin b) => [a] -> [b] -> [a]
pmod : {a, b} (fin a, fin b) => [a] -> [1 + b] -> [b]
pmult : {a, b} (fin a, fin b) => [a] -> [b] -> [max 1 (a + b) - 1]
```

Sequences

```
take : {front, back, elem} (fin front)
      => [front + back]elem -> [front]elem
drop : {front, back, elem} (fin front)
      => [front + back]elem -> [front]elem
tail : {a, b} [a+1]b -> [a]b
#     : {a, b, c} (fin a) => ([a]b, [c]b) -> [a+c]b
join  : {parts, each, a} (fin each)
      => [parts][each]a -> [parts * each]a
```

```

groupBy : {each, parts, elem} (fin each)
        => [parts * each]elem -> [parts][each]elem
splitBy : {parts, each, elem}
        => [parts * each]elem -> [parts][each]elem
reverse : {a, b} (fin a) => [a]b -> [a]b
@       : {a, b, c} ([a]b,[c]) -> b
!       : {a, b, c} (fin a) => ([a]b,[c]) -> b
@@      : {a, b, c, d} ([a]b,[c][d]) -> [c]b
!!     : {a, b, c, d} (fin a) => ([a]b,[c][d]) -> [c]b
width  : {a, b, c} (c >= width a) => [a]b -> [c]

```

Shifting, rotating

```

>>, <<, >>>, <<< : {a b c} (fin a,c >= lg2 a)
                  => ([a]b,[c]) -> [a]b

```

Miscellaneous

```

zero      : {a} a
transpose : {a, b, c} [a][b]c -> [b][a]c
min, max  : {a} (fin a) => ([a],[a]) -> [a]

```

Representing exceptions

```

error      : {a, b} [a][8] -> b
undefined  : {a} a

```

Appendix C

Enigma simulator

In this appendix we present the Cryptol code for the enigma machine in its entirety for reference purposes. Chapter 4 has a detailed discussion on how the enigma machine worked, and the construction of the Cryptol model below.

```
1 // Cryptol Enigma Simulator
2 // Copyright (c) 2010-2013, Galois Inc.
3 // www.cryptol.net
4 // You can freely use this source code for educational purposes.
5
6 // Helper synonyms:
7 // type Char      = [8]
8 module Enigma where
9
10 type Permutation = String 26
11
12 // Enigma components:
13 type Plugboard   = Permutation
14 type Rotor       = [26](Char, Bit)
15 type Reflector   = Permutation
16
17 // An enigma machine with n rotors:
18 type Enigma n    = { plugboard : Plugboard,
19                    rotors      : [n]Rotor,
20                    reflector   : Reflector
21                    }
22
23 // Check membership in a sequence:
24 elem : {a, b} (fin 0, fin a, Cmp b) => (b, [a]b) -> Bit
25 elem (x, xs) = matches ! 0
26   where matches = [False] # [ m || (x == e) | e <- xs
27                       | m <- matches
28                       ]
29
30 // Inverting a permutation lookup:
31 invSubst : (Permutation, Char) -> Char
32 invSubst (key, c) = candidates ! 0
33   where candidates = [0] # [ if c == k then a else p
34                             | k <- key
35                             | a <- ['A' .. 'Z']
36                             | p <- candidates
```

```

36         ]
37
38 // Constructing a rotor
39 mkRotor : {n} (fin n => (Permutation, String n) -> Rotor)
40 mkRotor (perm, notchLocations) = [ (p, elem (p, notchLocations))
41     | p <- perm
42     ]
43
44 // Action of a single rotor on a character
45 // Note that we encrypt and then rotate, if necessary
46 scramble : (Bit, Char, Rotor) -> (Bit, Char, Rotor)
47 scramble (rotate, c, rotor) = (notch, c', rotor')
48   where
49     (c', _) = rotor @ (c - 'A')
50     (_, notch) = rotor @ 0
51     rotor' = if rotate then rotor <<< 1 else rotor
52
53 // Connecting rotors in a sequence
54 joinRotors : {n} (fin n => ([n]Rotor, Char) -> ([n]Rotor, Char))
55 joinRotors (rotors, inputChar) = (rotors', outputChar)
56   where
57     initRotor = mkRotor (['A' .. 'Z'], [])
58     ncrs : [n+1](Bit, [8], Rotor)
59     ncrs = [(True, inputChar, initRotor)]
60         # [ scramble (notch, char, r)
61             | r <- rotors
62             | (notch, char, rotor') <- ncrs
63             ]
64     rotors' = tail [ r | (_, _, r) <- ncrs ]
65     (_, outputChar, _) = ncrs ! 0
66
67 // Following the signal through a single rotor, forward and backward
68 substFwd, substBwd : (Permutation, Char) -> Char
69 substFwd (perm, c) = perm @ (c - 'A')
70 substBwd (perm, c) = invSubst (perm, c)
71 // Route the signal back from the reflector, chase through rotors
72 backSignal : {n} (fin n => ([n]Rotor, Char) -> Char)
73 backSignal (rotors, inputChar) = cs ! 0
74   where cs = [inputChar] # [ substBwd ([ p | (p, _) <- r ], c)
75     | r <- reverse rotors
76     | c <- cs
77     ]
78
79 // The full enigma loop, from keyboard to lamps:
80 // The signal goes through the plugboard, rotors, and the reflector,
81 // then goes back through the sequence in reverse, out of the
82 // plugboard and to the lamps
83 enigmaLoop : {n} (fin n => (Plugboard, [n]Rotor, Reflector, Char) -> ([n]Rotor, Char))
84 enigmaLoop (pboard, rotors, refl, c0) = (rotors', c5)
85   where
86     c1 = substFwd (pboard, c0)
87     (rotors', c2) = joinRotors (rotors, c1)
88     c3 = substFwd (refl, c2)
89     c4 = backSignal (rotors, c3)
90     c5 = substBwd (pboard, c4)
91
92 // Construct a machine out of parts
93 mkEnigma : {n} (Plugboard, [n]Rotor, Reflector, [n]Char) -> Enigma n

```

```

94 mkEnigma (pboard, rs, refl, startingPositions) =
95   { plugboard = pboard
96     , rotors   = [ r <<< (s - 'A')
97                   | r <- rs
98                   | s <- startingPositions
99                   ]
100   , reflector = refl
101   }
102
103 // Encryption/Decryption
104 enigma : {n, m} (fin n, fin m) => (Enigma n, String m) -> String m
105 enigma (m, pt) = tail [ c | (_, c) <- rcs ]
106   where rcs = [(m.rotors, '*')] #
107             [ enigmaLoop (m.plugboard, r, m.reflector, c)
108               | c <- pt
109               | (r, _) <- rcs
110             ]
111
112 // Decryption is the same as encryption:
113 // dEnigma : {n, m} (fin n, fin m) => (Enigma n, String m) -> String m
114 dEnigma = enigma
115
116
117 // Build an example enigma machine:
118 plugboard : Plugboard
119 plugboard = "HBGDEFCAIJKOWNLPXRSVYTMQUZ"
120
121 rotor1, rotor2, rotor3 : Rotor
122 rotor1 = mkRotor ("RJICAWVQZODLUPYFEHXSMTKNGB", "IO")
123 rotor2 = mkRotor ("DWYOLETKNVQPHURZJMSFIGXCBA", "B")
124 rotor3 = mkRotor ("FGKMAJWUOVNRYIZETDPSHBLQX", "CK")
125
126 reflector : Reflector
127 reflector = "FEIPBATSCYVUWZQDOXHGLKMRJN"
128
129 modelEnigma : Enigma 3
130 modelEnigma = mkEnigma (plugboard, [rotor1, rotor2, rotor3], reflector, "GCR")
131
132 /* Example run:
133
134 cryptol> :set ascii=on
135 cryptol> enigma (modelEnigma, "ENIGMAWASAREALLYCOOLMACHINE")
136 UPEKTBSDROBVTUJGNCEHHGBXGTF
137 cryptol> dEnigma (modelEnigma, "UPEKTBSDROBVTUJGNCEHHGBXGTF")
138 ENIGMAWASAREALLYCOOLMACHINE
139 */
140
141 all : {a, n} (fin n) => (a->Bit) -> [n]a -> Bit
142 all fn xs = folds ! 0 where
143   folds = [True] # [ fn x && p | x <- xs
144                       | p <- folds]
145 checkReflectorFwdBwd : Reflector -> Bit
146 checkReflectorFwdBwd refl = all check ['A' .. 'Z']
147   where check c = substFwd (refl, c) == substBwd (refl, c)
148

```


Appendix D

AES in Cryptol

In this appendix we present the Cryptol code for the AES in its entirety for reference purposes. Chapter 6 has a detailed discussion on how AES works, and the construction of the Cryptol model below.

In the below code, simply set `Nk` to be 4 for AES128, 6 for AES192, and 8 for AES256 on line 15. No other modifications are required for obtaining these AES variants. Note that we have rearranged the code given in Chapter 6 below for ease of reading.

```
1 // Cryptol AES Implementation
2 // Copyright (c) 2010-2013, Galois Inc.
3 // www.cryptol.net
4 // You can freely use this source code for educational purposes.
5
6 // This is a fairly close implementation of the FIPS-197 standard:
7 // http://csrc.nist.gov/publications/fips/fips197/fips-197.pdf
8
9 // Nk: Number of blocks in the key
10 // Must be one of 4 (AES128), 6 (AES192), or 8 (AES256)
11 // Aside from this line, no other code below needs to change for
12 // implementing AES128, AES192, or AES256
13 module AES where
14
15 type AES128 = 4
16 type AES192 = 6
17 type AES256 = 8
18
19 type Nk = AES128
20
21 // For Cryptol 2.x | x > 0
22 // NkValid: 'Nk -> Bit
23 // property NkValid k = (k == 'AES128) || (k == 'AES192) || (k == 'AES256)
24
25 // Number of blocks and Number of rounds
26 type Nb = 4
27 type Nr = 6 + Nk
28
29 type AESKeySize = (Nk*32)
30
31 // Helper type definitions
32 type GF28 = [8]
33 type State = [4][Nb]GF28
```

```

34 type RoundKey = State
35 type KeySchedule = (RoundKey, [Nr-1]RoundKey, RoundKey)
36 // GF28 operations
37 gf28Add : {n} (fin n) => [n]GF28 -> GF28
38 gf28Add ps = sums ! 0
39   where sums = [zero] # [ p ^ s | p <- ps | s <- sums ]
40
41 irreducible = <| x^^8 + x^^4 + x^^3 + x + 1 |>
42
43 gf28Mult : (GF28, GF28) -> GF28
44 gf28Mult (x, y) = pmod(pmult x y) irreducible
45
46 gf28Pow : (GF28, [8]) -> GF28
47 gf28Pow (n, k) = pow k
48   where sq x = gf28Mult (x, x)
49         odd x = x ! 0
50         pow i = if i == 0 then 1
51               else if odd i
52                   then gf28Mult(n, sq (pow (i >> 1)))
53                   else sq (pow (i >> 1))
54
55 gf28Inverse : GF28 -> GF28
56 gf28Inverse x = gf28Pow (x, 254)
57
58 gf28DotProduct : {n} (fin n) => ([n]GF28, [n]GF28) -> GF28
59 gf28DotProduct (xs, ys) = gf28Add [ gf28Mult (x, y) | x <- xs
60                                   | y <- ys ]
61
62 gf28VectorMult : {n, m} (fin n) => ([n]GF28, [m][n]GF28) -> [m]GF28
63 gf28VectorMult (v, ms) = [ gf28DotProduct(v, m) | m <- ms ]
64
65 gf28MatrixMult : {n, m, k} (fin m) => ([n][m]GF28, [m][k]GF28) -> [n][k]GF28
66 gf28MatrixMult (xss, yss) = [ gf28VectorMult(xs, yss') | xs <- xss ]
67   where yss' = transpose yss
68
69 // The affine transform and its inverse
70 xformByte : GF28 -> GF28
71 xformByte b = gf28Add [b, (b >>> 4), (b >>> 5), (b >>> 6), (b >>> 7), c]
72   where c = 0x63
73
74 xformByte' : GF28 -> GF28
75 xformByte' b = gf28Add [(b >>> 2), (b >>> 5), (b >>> 7), d] where d = 0x05
76 // The SubBytes transform and its inverse
77 SubByte : GF28 -> GF28
78 SubByte b = xformByte (gf28Inverse b)
79
80 SubByte' : GF28 -> GF28
81 SubByte' b = sbox@b
82
83 SubBytes : State -> State
84 SubBytes state = [ [ SubByte' b | b <- row ] | row <- state ]
85
86
87 InvSubByte : GF28 -> GF28
88 InvSubByte b = gf28Inverse (xformByte' b)
89
90 InvSubBytes : State -> State
91 InvSubBytes state = [ [ InvSubByte b | b <- row ] | row <- state ]

```

```

92
93 // The ShiftRows transform and its inverse
94 ShiftRows : State -> State
95 ShiftRows state = [ row <<< shiftAmount | row <- state
96                               | shiftAmount <- [0 .. 3]
97                               ]
98
99 InvShiftRows : State -> State
100 InvShiftRows state = [ row >>> shiftAmount | row <- state
101                               | shiftAmount <- [0 .. 3]
102                               ]
103
104 // The MixColumns transform and its inverse
105 MixColumns : State -> State
106 MixColumns state = gf28MatrixMult (m, state)
107     where m = [[2, 3, 1, 1],
108               [1, 2, 3, 1],
109               [1, 1, 2, 3],
110               [3, 1, 1, 2]]
111
112 InvMixColumns : State -> State
113 InvMixColumns state = gf28MatrixMult (m, state)
114     where m = [[0x0e, 0x0b, 0x0d, 0x09],
115               [0x09, 0x0e, 0x0b, 0x0d],
116               [0x0d, 0x09, 0x0e, 0x0b],
117               [0x0b, 0x0d, 0x09, 0x0e]]
118
119 // The AddRoundKey transform
120 AddRoundKey : (RoundKey, State) -> State
121 AddRoundKey (rk, s) = rk ^ s
122 // Key expansion
123 Rcon : [8] -> [4]GF28
124 Rcon i = [(gf28Pow (<| x |>, i-1)), 0, 0, 0]
125
126 SubWord : [4]GF28 -> [4]GF28
127 SubWord bs = [ SubByte b | b <- bs ]
128
129 RotWord : [4]GF28 -> [4]GF28
130 RotWord [a0, a1, a2, a3] = [a1, a2, a3, a0]
131
132 NextWord : ([8],[4][8],[4][8]) -> [4][8]
133 NextWord(i, prev, old) = old ^ mask
134     where mask = if i % 'Nk == 0
135                   then SubWord(RotWord(prev)) ^ Rcon (i / 'Nk)
136                   else if ('Nk > 6) && (i % 'Nk == 4)
137                         then SubWord(prev)
138                         else prev
139
140
141 ExpandKeyForever : [Nk][4][8] -> [inf]RoundKey
142 ExpandKeyForever seed = [ transpose g | g <- groupBy' {4} (keyWS seed) ]
143
144 keyWS : [Nk][4][8] -> [inf][4][8]
145 keyWS seed = xs
146     where xs = seed # [ NextWord(i, prev, old)
147                       | i <- [ 'Nk ... ]
148                       | prev <- drop' {Nk-1} xs
149                       | old <- xs

```

```

150         ]
151
152 checkKey = take'{16} (drop'{8} (keyWS ["abcd", "defg", "1234", "5678"]))
153 checkKey2 = [transpose g | g <- groupBy'{4}checkKey]
154
155 ExpandKey : [AESKeySize] -> KeySchedule
156 ExpandKey key = (keys @ 0, keys @ [1 .. (Nr - 1)], keys @ 'Nr)
157   where seed : [Nk][4][8]
158         seed = split (split key)
159         keys = ExpandKeyForever seed
160
161 fromKS : KeySchedule -> [Nr+1][4][32]
162 fromKS (f, ms, l) = [ formKeyWords (transpose k) | k <- [f] # ms # [l] ]
163   where formKeyWords bbs = [ join bs | bs <- bbs ]
164
165 // AES rounds and inverses
166 AESRound : (RoundKey, State) -> State
167 AESRound (rk, s) = AddRoundKey (rk, MixColumns (ShiftRows (SubBytes s)))
168
169 AESFinalRound : (RoundKey, State) -> State
170 AESFinalRound (rk, s) = AddRoundKey (rk, ShiftRows (SubBytes s))
171
172 AESInvRound : (RoundKey, State) -> State
173 AESInvRound (rk, s) =
174   InvMixColumns (AddRoundKey (rk, InvSubBytes (InvShiftRows s)))
175 AESFinalInvRound : (RoundKey, State) -> State
176 AESFinalInvRound (rk, s) = AddRoundKey (rk, InvSubBytes (InvShiftRows s))
177
178 // Converting a 128 bit message to a State and back
179 msgToState : [128] -> State
180 msgToState msg = transpose (split (split msg))
181
182 stateToMsg : State -> [128]
183 stateToMsg st = join (join (transpose st))
184
185 // AES Encryption
186 aesEncrypt : ([128], [AESKeySize]) -> [128]
187 aesEncrypt (pt, key) = stateToMsg (AESFinalRound (kFinal, rounds ! 0))
188   where (kInit, ks, kFinal) = ExpandKey key
189         state0 = AddRoundKey(kInit, msgToState pt)
190         rounds = [state0] # [ AESRound (rk, s) | rk <- ks
191                               | s <- rounds
192                               ]
193
194 // AES Decryption
195 aesDecrypt : ([128], [AESKeySize]) -> [128]
196 aesDecrypt (ct, key) = stateToMsg (AESFinalInvRound (kFinal, rounds ! 0))
197   where (kFinal, ks, kInit) = ExpandKey key
198         state0 = AddRoundKey(kInit, msgToState ct)
199         rounds = [state0] # [ AESInvRound (rk, s)
200                               | rk <- reverse ks
201                               | s <- rounds
202                               ]
203
204 sbox : [256]GF28
205 sbox = [
206   0x63, 0x7c, 0x77, 0x7b, 0xf2, 0x6b, 0x6f, 0xc5, 0x30, 0x01, 0x67,
207   0x2b, 0xfe, 0xd7, 0xab, 0x76, 0xca, 0x82, 0xc9, 0x7d, 0xfa, 0x59,

```

```

208 0x47, 0xf0, 0xad, 0xd4, 0xa2, 0xaf, 0x9c, 0xa4, 0x72, 0xc0, 0xb7,
209 0xfd, 0x93, 0x26, 0x36, 0x3f, 0xf7, 0xcc, 0x34, 0xa5, 0xe5, 0xf1,
210 0x71, 0xd8, 0x31, 0x15, 0x04, 0xc7, 0x23, 0xc3, 0x18, 0x96, 0x05,
211 0x9a, 0x07, 0x12, 0x80, 0xe2, 0xeb, 0x27, 0xb2, 0x75, 0x09, 0x83,
212 0x2c, 0x1a, 0x1b, 0x6e, 0x5a, 0xa0, 0x52, 0x3b, 0xd6, 0xb3, 0x29,
213 0xe3, 0x2f, 0x84, 0x53, 0xd1, 0x00, 0xed, 0x20, 0xfc, 0xb1, 0x5b,
214 0x6a, 0xcb, 0xbe, 0x39, 0x4a, 0x4c, 0x58, 0xcf, 0xd0, 0xef, 0xaa,
215 0xfb, 0x43, 0x4d, 0x33, 0x85, 0x45, 0xf9, 0x02, 0x7f, 0x50, 0x3c,
216 0x9f, 0xa8, 0x51, 0xa3, 0x40, 0x8f, 0x92, 0x9d, 0x38, 0xf5, 0xbc,
217 0xb6, 0xda, 0x21, 0x10, 0xff, 0xf3, 0xd2, 0xcd, 0x0c, 0x13, 0xec,
218 0x5f, 0x97, 0x44, 0x17, 0xc4, 0xa7, 0x7e, 0x3d, 0x64, 0x5d, 0x19,
219 0x73, 0x60, 0x81, 0x4f, 0xdc, 0x22, 0x2a, 0x90, 0x88, 0x46, 0xee,
220 0xb8, 0x14, 0xde, 0x5e, 0x0b, 0xdb, 0xe0, 0x32, 0x3a, 0x0a, 0x49,
221 0x06, 0x24, 0x5c, 0xc2, 0xd3, 0xac, 0x62, 0x91, 0x95, 0xe4, 0x79,
222 0xe7, 0xc8, 0x37, 0x6d, 0x8d, 0xd5, 0x4e, 0xa9, 0x6c, 0x56, 0xf4,
223 0xea, 0x65, 0x7a, 0xae, 0x08, 0xba, 0x78, 0x25, 0x2e, 0x1c, 0xa6,
224 0xb4, 0xc6, 0xe8, 0xdd, 0x74, 0x1f, 0x4b, 0xbd, 0x8b, 0x8a, 0x70,
225 0x3e, 0xb5, 0x66, 0x48, 0x03, 0xf6, 0x0e, 0x61, 0x35, 0x57, 0xb9,
226 0x86, 0xc1, 0x1d, 0x9e, 0xe1, 0xf8, 0x98, 0x11, 0x69, 0xd9, 0x8e,
227 0x94, 0x9b, 0x1e, 0x87, 0xe9, 0xce, 0x55, 0x28, 0xdf, 0x8c, 0xa1,
228 0x89, 0x0d, 0xbf, 0xe6, 0x42, 0x68, 0x41, 0x99, 0x2d, 0x0f, 0xb0,
229 0x54, 0xbb, 0x16]
230
231 // Test runs:
232
233 // cryptol> aesEncrypt (0x3243f6a8885a308d313198a2e0370734, \
234 //                      0x2b7e151628aed2a6abf7158809cf4f3c)
235 // 0x3925841d02dc09fbdcl18597196a0b32
236 // cryptol> aesEncrypt (0x00112233445566778899aabbccddeeff, \
237 //                      0x000102030405060708090a0b0c0d0e0f)
238 // 0x69c4e0d86a7b0430d8cdb78070b4c55a
239 property AESCorrect msg key = aesDecrypt (aesEncrypt (msg, key), key) == msg

```


Appendix E

Cryptol Syntax

E.1 Layout

Groups of declarations are organized based on indentation. Declarations with the same indentation belong to the same group. Lines of text that are indented more than the beginning of a declaration belong to that declaration, while lines of text that are indented less terminate a group of declaration. Groups of declarations appear at the top level of a Cryptol file, and inside `where` blocks in expressions. For example, consider the following declaration group

```
f x = x + y + z
  where
    y = x * x
    z = x + y

g y = y
```

This group has two declarations, one for `f` and one for `g`. All the lines between `f` and `g` that are indented more than `f` belong to `f`.

This example also illustrates how groups of declarations may be nested within each other. For example, the `where` expression in the definition of `f` starts another group of declarations, containing `y` and `z`. This group ends just before `g`, because `g` is indented less than `y` and `z`.

E.2 Comments

Cryptol supports block comments, which start with `/*` and end with `*/`, and line comments, which start with `//` and terminate at the end of the line. Block comments may be nested arbitrarily.

Examples:

```
/* This is a block comment */
// This is a line comment
/* This is a /* Nested */ block comment */
```

E.3 Identifiers

Cryptol identifiers consist of one or more characters. The first character must be either an English letter or underscore (`_`). The following characters may be an English letter, a decimal digit, underscore (`_`), or a prime (`'`). Some identifiers have special meaning in the language, so they may not be used in programmer-defined names (see Keywords).

Examples:

```
name    name1    name'    longer_name
Name    Name2    Name''   longerName
```

E.4 Keywords and Built-in Operators

The following identifiers have special meanings in Cryptol, and may not be used for programmer defined names:

```
Arith  Inf      extern  inf      module   then
Bit     True     fin     lg2     newtype  type
Cmp     else     if      max     pragma   where
False   export   import  min     property width
```

The following table contains Cryptol's operators and their associativity with lowest precedence operators first, and highest precedence last.

Operator	Associativity
<code> </code>	left
<code>^</code>	left
<code>&&</code>	left
<code>-> (types)</code>	right
<code>!= ==</code>	not associative
<code>> < <= >=</code>	not associative
<code>#</code>	right
<code>>> << >>> <<<</code>	left
<code>+ -</code>	left
<code>* / %</code>	left
<code>^^</code>	right
<code>! !! @ @@</code>	left
(unary) <code>- ~</code>	right

Table E.1: Operator precedences.

E.5 Numeric Literals

Numeric literals may be written in binary, octal, decimal, or hexadecimal notation. The base of a literal is determined by its prefix: `0b` for binary, `0o` for octal, no special prefix for decimal, and `0x` for hexadecimal.

Examples:

```
254           // Decimal literal
0254          // Decimal literal
0b11111110   // Binary literal
0o376         // Octal literal
0xFE          // Hexadecimal literal
0xfe         // Hexadecimal literal
```

Numeric literals represent finite bit sequences (i.e., they have type `[n]`). Using binary, octal, and hexadecimal notation results in bit sequences of a fixed length, depending on the number of digits in the literal. Decimal literals are overloaded, and so the length of the sequence is inferred from context in which the literal is used. Examples:

```
0b1010       // : [4], 1 * number of digits
0o1234       // : [12], 3 * number of digits
0x1234       // : [16], 4 * number of digits

10           // : {n}. (fin n, n >= 4) => [n]
             // (need at least 4 bits)

0            // : {n}. (fin n) => [n]
```

E.6 Bits

The type `Bit` has two inhabitants: `True` and `False`. These values may be combined using various logical operators, or constructed as results of comparisons.

Operator	Associativity	Description
<code> </code>	left	Logical or
<code>^</code>	left	Exclusive-or
<code>&&</code>	left	Logical and
<code>!= ==</code>	none	Not equals, equals
<code>> < <= >=</code>	none	Comparisons
<code>~</code>	right	Logical negation

Table E.2: Bit operations.

E.7 If Then Else with Multiway

If `then else` has been extended to support multi-way conditionals. Examples:

```
x = if y % 2 == 0 then 22 else 33
```

```
x = if y % 2 == 0 then 1
    | y % 3 == 0 then 2
    | y % 5 == 0 then 3
    else 7
```

E.8 Tuples and Records

Tuples and records are used for packaging multiples values together. Tuples are enclosed in parenthesis, while records are enclosed in braces. The components of both tuples and records are separated by commas. The components of tuples are expressions, while the components of records are a label and a value separated by an equal sign. Examples:

```
(1,2,3)           // A tuple with 3 component
()                // A tuple with no components

{ x = 1, y = 2 }  // A record with two fields, 'x' and 'y'
{}               // A record with no fields
```

The components of tuples are identified by position, while the components of records are identified by their label, and so the ordering of record components is not important. Examples:

```
(1,2) == (1,2)           // True
(1,2) == (2,1)           // False

{ x = 1, y = 2 } == { x = 1, y = 2 } // True
{ x = 1, y = 2 } == { y = 2, x = 1 } // True
```

The components of a record or a tuple may be accessed in two ways: via pattern matching or by using explicit component selectors. Explicit component selectors are written as follows:

```
(15, 20).0 == 15
(15, 20).1 == 20
```

```
{ x = 15, y = 20 }.x == 15
```

Explicit record selectors may be used only if the program contains sufficient type information to determine the shape of the tuple or record. For example:

```
type T = { sign :: Bit, number :: [15] }
```

```
// Valid definition:
```

```
// the type of the record is known.
isPositive : T -> Bit
isPositive x = x.sign

// Invalid definition:
// insufficient type information.
badDef x = x.f
```

The components of a tuple or a record may also be accessed by using pattern matching. Patterns for tuples and records mirror the syntax for constructing values: tuple patterns use parenthesis, while record patterns use braces. Examples:

```
getFst (x,_) = x

distance2 { x = xPos, y = yPos } = xPos ^^ 2 + yPos ^^ 2

f x = fst + snd where
```

E.9 Sequences

A sequence is a fixed-length collection of elements of the same type. The type of a finite sequence of length n , with elements of type a is $[n] a$. Often, a finite sequence of bits, $[n] \text{ Bit}$, is called a *word*. We may abbreviate the type $[n] \text{ Bit}$ as $[n]$. An infinite sequence with elements of type a has type $[\text{inf}] a$, and $[\text{inf}]$ is an infinite stream of bits.

```
[e1,e2,e3]           // A sequence with three elements

[t .. ]             // Sequence enumerations
[t1, t2 .. ]        // Step by t2 - t1
[t1 .. t3 ]
[t1, t2 .. t3 ]
[e1 ... ]           // Infinite sequence starting at e1
[e1, e2 ... ]       // Infinite sequence stepping by e2-e1

[ e | p11 <- e11, p12 <- e12 // Sequence comprehensions
  | p21 <- e21, p22 <- e22 ]
```

Note: the bounds in finite unbounded (those with `..`) sequences are type expressions, while the bounds in bounded-finite and infinite sequences are value expressions.

Operator	Description
#	Sequence concatenation
>> <<	Shift (right,left)
>>> <<<	Rotate (right,left)

Operator	Description
<code>e !</code>	Access elements (front,back)
<code>@@ !!</code>	Access sub-sequence (front,back)

Table E.3: Sequence operations.

There are also lifted point-wise operations.

```
[p1, p2, p3, p4]      // Sequence pattern
p1 # p2              // Split sequence pattern
```

E.10 Functions

```
\p1 p2 -> e          // Lambda expression
f p1 p2 = e          // Function definition
```

E.11 Local Declarations

```
e where ds
```

E.12 Explicit Type Instantiation

If `f` is a polymorphic value with type:

```
f : { tyParam }
f '{ tyParam = t }
```

E.13 Demoting Numeric Types to Values

The value corresponding to a numeric type may be accessed using the following notation:

```
'{t}
```

Here `t` should be a type expression with numeric kind. The resulting expression is a finite word, which is sufficiently large to accommodate the value of the type:

```
'{t} :: {w >= width t}. [w]
```

E.14 Explicit Type Annotations

Explicit type annotations may be added on expressions, patterns, and in argument definitions.

```
e : t
```

```
p : t
```

```
f (x : t) = ...
```

E.15 Type Signatures

```
f,g : {a,b} (fin a) => [a] b
```

E.16 Type Synonym Declarations

```
type T a b = [a] b
```


Appendix F

The Cryptol Grammar

Glossary

AES The Advanced Encryption Standard [13], 61

Cipherkey The key used in a particular encryption/decryption task, 37

Ciphertext The result of encrypting a plaintext message, “unreadable” unless the key is known, 37

Fibonacci numbers The sequence 0, 1, 1, 2, 3, 5, . . . After the elements 0 and 1, each consecutive element is the sum of the two previous numbers [20], 31

NIST National Institute of Standards and Technology. The institution in charge of standardizing cryptographic algorithms (amongst many other things) in USA., 61

Plaintext A “readable” message that we would like to encrypt, the message in the clear, 37

SAT Solver An automated tool for solving boolean satisfiability problems. Cryptol uses SAT/SMT solvers in order to provide its high-assurance capabilities, 54, 57, 59

SMT Solver Satisfiability Modulo Theories: An automated tool for establishing satisfiability problems with respect to certain theories. One of the theories of interest to Cryptol is that of bit-vectors, as it provides a natural medium for translating Cryptol’s bit-precise theorems, 54, 57, 59

Bibliography

(Each entry is followed by a list of page numbers on which the citation appears. All cited URLs, unless otherwise stated, were last accessed in November 2010.)

- [1] Richard Bird. *Introduction to Functional Programming using Haskell*. Prentice Hall Europe, London, second edition, 1998. 29
- [2] Andy Carlson. Simulating the enigma cypher machine. http://homepages.tesco.net/~andycarlson/enigma/simulating_enigma.html, see the section “Wheel Turnover and The Anomaly”. 46, 102
- [3] Koen Claessen and John Hughes. QuickCheck: A lightweight tool for random testing of Haskell programs. In *Proc. of International Conference on Functional Programming (ICFP)*. ACM SIGPLAN, 2000. 57
- [4] Cryptol web site. <http://www.cryptol.net>. 3
- [5] Joan Daemen and Vincent Rijmen. *The Design of Rijndael: AES - The Advanced Encryption Standard*. Springer, 2002. 61
- [6] Levent Erkök, Magnus Carlsson, and Adam Wick. Hardware/software co-verification of cryptographic algorithms using Cryptol. In *Formal Methods in Computer Aided Design, FMCAD'09, Austin, TX, USA*, pages 188–191. IEEE, November 2009. 22
- [7] Levent Erkök and John Matthews. Pragmatic equivalence and safety checking in Cryptol. In *Programming Languages meets Program Verification, PLPV'09, Savannah, Georgia, USA*, pages 73–81. ACM Press, January 2009. 22, 25, 54, 55
- [8] J. Roger Hindley. *Basic Simple Type Theory*, volume 42. Cambridge University Press, Cambridge, UK, 1997. 22
- [9] Brian W. Kernighan and Dennis M. Richie. *The C Programming Language*. Prentice Hall, second edition, 1998. 32
- [10] Daniel Kroening and Ofer Strichman. *Decision Procedures: An Algorithmic Point of View*. Springer, 2008. 108
- [11] J. R. Lewis and B. Martin. Cryptol: high assurance, retargetable crypto development and validation. In *Military Communications Conference 2003*, volume 2, pages 820–825. IEEE, October 2003. 26

- [12] Robin Milner, Mads Tofte, and David MacQueen. *The Definition of Standard ML*. MIT Press, Cambridge, MA, USA, 1997. 23
- [13] National Institute of Standards and Technology, NIST. Announcing the AES. <http://csrc.nist.gov/publications/fips/fips197/fips-197.pdf>, November 2001. FIPS Publication 197. 61, 62, 64, 66, 67, 68, 69, 70, 71, 72, 73, 74, 75, 137
- [14] Simon L. Peyton Jones and John Hughes. (Editors.) Report on the programming language Haskell 98, a non-strict purely-functional programming language. URL: www.haskell.org/onlinereport, February 1999. 23, 27, 32
- [15] Simon Singh. *The code book: the evolution of secrecy from Mary, Queen of Scots, to quantum cryptography*. Doubleday, New York, NY, USA, 1999. 37, 43, 47
- [16] CVC4 web site. <http://cvc4.cs.nyu.edu/web/>. 55
- [17] Z3 web site. <http://z3.codeplex.com/>. 55
- [18] Wikipedia. Enigma machine — wikipedia, the free encyclopedia. http://en.wikipedia.org/w/index.php?title=Enigma_machine&oldid=392040616, 2010. [Online; accessed 21-October-2010]. 43
- [19] Wikipedia. Enigma rotor details — wikipedia, the free encyclopedia. http://en.wikipedia.org/w/index.php?title=Enigma_rotor_details&oldid=389862975, 2010. [Online; accessed 25-October-2010]. 50
- [20] Wikipedia. Fibonacci number — wikipedia, the free encyclopedia. http://en.wikipedia.org/w/index.php?title=Fibonacci_number&oldid=390711214, 2010. [Online; accessed 15-October-2010]. 31, 137
- [21] Wikipedia. Finite field arithmetic — wikipedia, the free encyclopedia, 2010. [Online; accessed 10-November-2010]. 62
- [22] Wikipedia. Polynomial long division — wikipedia, the free encyclopedia, 2010. [Online; accessed 10-November-2010]. 64
- [23] Wikipedia. Scytale — wikipedia, the free encyclopedia. <http://en.wikipedia.org/w/index.php?title=Scytale&oldid=391405769>, 2010. [Online; accessed 19-October-2010]. 41
- [24] Wikipedia. Substitution cipher — wikipedia, the free encyclopedia. http://en.wikipedia.org/w/index.php?title=Substitution_cipher&oldid=389116332, 2010. [Online; accessed 20-October-2010]. 40
- [25] Wikipedia. Vigenere cipher — wikipedia, the free encyclopedia. http://en.wikipedia.org/w/index.php?title=Vigenere_cipher&oldid=389811286, 2010. [Online; accessed 19-October-2010]. 38
- [26] Yices web site. <http://yices.csl.sri.com/>. 55

Index

GF(2⁸), *see* galois field
λ-expression, 27, 28, 56, 58, 104

&, (and), 9
#, (append), 13, 15, 52, 82, 102
~, (complement), 9, 19, 53
/, (divide), 20, 56, 57, 86
==, (equal), 20, 85, 90
**, (exponentiate), 20, 51
>, (greater than), 20
>=, (greater or equal), 20
<, (less than), 20
<=, (less or equal), 20
-, (subtract), 20, 51, 90
%, (modulus), 20, 56, 57, 86
|, (or), 9, 91
+, (add), 20, 26, 51
!, (reverse select), 13, 29, 30, 46, 72, 90, 96
!!, (reverse permutation), 13
«<, (rotate left), 16, 17, 38, 45, 66, 69, 93, 94
»>, (rotate right), 16, 17, 93, 94
@, (select), 13, 17, 40, 45, 70
«, (shift left), 16, 17
», (shift right), 16, 17
*, (multiply), 20, 51, 56, 57
-, (unary minus), 20
_, (underscore), 45
^, (xor), 9, 63, 71
!=, (not-equal), 20, 85
AES, 61, 62, 64–75, 121
 InvMixColumns, 74
 InvShiftRows, 74
 InvSubBytes, 73
 MixColumns, 67, 74
 ShiftRows, 66, 74
 state, 62
 SubBytes, 64, 74

all, 28, 57, 89, 101, 104
ambiguous constraints, *see* type, ambiguous
any, 28, 89
atbash, 39, 96

bit, 9, 22

Caesar's cipher, 38, 39, 57, 96
case sensitivity, 4, 10
characters, 17
cipherkey, 37, 41
ciphertext, 37, 41
command line

 completion, 5
 line continuation, 4, 80, 83

commands

 :b (browse), 6
 :check, 57, 58, 67, 75, 108
 :e (edit), 6
 :? (help), 6
 :i (info), 52, 102
 :l (load), 6, 26
 :m (module load), 6
 :p (print), 18
 :prove, 54, 56–58, 108
 :q (quit), 6
 :r (reload), 6, 26
 :sat, 58, 59
 ! (shell), 6
 :t (type), 5

comments, 4

defaulting, *see* type, defaulting
drop, 14, 15, 17, 70, 84, 92

elem, 31, 44, 57, 58, 104, 111
endianness, 16, 17, 83
Enigma machine, 43, 117

- plugboard, 43
- reflector, 47
- rotor, 44
- False, 9, 19, 22
- fin, *see* type, fin
- floating point, 10
- fold, 29, 31, 41, 46, 48, 49, 63, 97
- function application, 26, 27
- Galois field, 62, 64, 68, 74
- group, 14, 15, 17, 24, 82, 84, 87
- import directive, 5
- inference, *see* type, inference
- join, 14, 15, 41, 42, 72, 82
- known plaintext attack, 39
- lambda expression, *see* λ -expression
- lg2, 20, 26, 85
- literate programming, 5
- max, 20, 29, 88
- min, 20, 88
- modular arithmetic, 20, 21, 52, 55, 86, 91, 92
- module system, 5
- monomorphism, *see* type, monomorphism
- overflow, 20, 91
- overloading, *see* type, overloading
- pattern matching, 45, 46
- pdiv, *see* polynomial, division
- plaintext, 37, 38, 41
- pmod, *see* polynomial, modulus
- pmult, *see* polynomial, multiplication
- polymorphism, *see* type, polymorphism
- polynomial, 62
 - addition, 63
 - division, 64
 - irreducible, 64
 - modulus, 64
 - multiplication, 63
 - subtraction, 63
- predicates, *see* type, predicates
- private qualifier, 5
- project, 100
- properties, 51
 - Q.E.D., 55
 - contradiction, 55
 - completeness, 57
 - conditional, 56, 57, 111
 - counter-example, 55, 57
 - function correspondence, 52, 53
 - polymorphic validity, 54
 - proving, 55
- record, 18, 23, 48
- recurrences, 28, 29
- recursion, 28
- reverse, 27, 48, 52, 72, 89, 102
- satisfiability checking, 58
- scytale, 41, 42, 98
- sequence, 11, 22
 - arithmetic lifting, 86
 - cartesian, 12, 13
 - comprehensions, 12, 13, 15, 29, 31, 46, 82, 86, 97
 - enumerations, 12, 21, 38, 86
 - finite, 14
 - infinite, 14, 21, 22, 86
 - nested, 13
 - parallel, 13
- settings
 - a, (ASCII printing), 4, 18, 19, 38, 77
 - base, (output base), 10, 16, 77, 83
 - editor, (file editor), 6
 - quickCheckCount, 58, 105
 - sbv, 107
 - t, (printing types), 16
- shift cipher, 37
- signature, *see* type, signature
- split, 15, 17, 41, 42, 72, 82, 84, 98
- splitBy, 14, 15, 82
- stream equations, 31
- streams, 14, 28
- string, 17
- substitution cipher, 37, 40
 - monoalphabetic, 40
 - period, 50, 102
 - polyalphabetic, 40, 43, 44, 50, 99
 - polygraphic, 40

symmetric key, 61

tail, 23, 27, 46, 50, 89

take, 14, 15, 17, 25, 83, 84

transpose, 14, 15, 41, 70, 72, 82, 111

transposition cipher, 37, 41

True, 9, 19, 22, 89

tuple, 11, 18, 22

type

 fin, 25, 27, 53, 87, 90

 inf, 22, 25

 ambiguous, 98

 annotation, 56, 104

 defaulting, 16

 inference, 5, 22

 inline argument types, 35

 monomorphism, 22, 54

 overloading, 23

 polymorphism, 5, 23, 24

 positional arguments, 34

 predicates, 25, 28, 39, 87, 90

 signature, 5, 15, 16, 41, 51, 87, 89, 98,
 104

 type classes, 33

 type context, 35

 type variables, 33

 undecidable, 25, 87

type synonym, 32

 Bool, 33

 String, 33, 38

 Word, 33

undecidable, *see* type, undecidable

underflow, 20

Vigenère cipher, 38, 96

where clause, 27, 28, 38, 56, 104

while loop, 30

width, 27, 54, 89, 103

wildcard, *see* _ (underscore)

word, 10, 16, 22

 arbitrary precision, 10

Yices, 55

zero, 19, 28, 53, 85, 89, 103