

SCREAM

Rob King

Introduction

- The Problem
- The Base64 Algorithm
- Regular Expressions

The Algorithm

- Ways of Solving the Problem
- Encoding Operations

Performance

- Expression Optimization
- Performance Analysis

Implementation and Usage

- Common Use Cases
- Caveats

Summary

Static Compilation of Regular Expressions for Analysis and Modification

Rob King

DVLabs
TippingPoint Technologies

March 25, 2010

SCREAM

Rob King

Introduction

The Problem

The Base64

Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression

Optimization

Performance

Analysis

Implementation
and Usage

Common Use Cases

Caveats

Summary

1

Introduction

- The Problem
- The Base64 Algorithm
- Regular Expressions

2

The Algorithm

- Ways of Solving the Problem
- Encoding Operations

3

Performance

- Expression Optimization
- Performance Analysis

4

Implementation and Usage

- Common Use Cases
- Caveats

5

Summary

THE PROBLEM

In which we discover purpose of the whole thing...

SCREAM

Rob King

Introduction

The Problem

The Base64

Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression

Optimization

Performance

Analysis

Implementation and Usage

Common Use Cases

Caveats

Summary

- This talk is about inspecting streams of data for interesting patterns, even when that stream of data has been encoded.
- We focus on the Base64 encoding scheme, and discuss a tool that can be used when dealing with Base64.
- However, most portions of the algorithm are applicable to other position-dependent bitwise block encodings (and, potentially, self-synchronizing encodings).
- We also want to talk about how Erlang made development of the tool much easier.
 - This tool also provided an interesting “back door” to get Erlang accepted in DVLabs.

THE PROBLEM

In which we discover purpose of the whole thing...

SCREAM

Rob King

Introduction

The Problem

The Base64

Algorithm

Regular Expressions

The Algorithm

Ways of Solving the Problem

Encoding Operations

Performance

Expression

Optimization

Performance

Analysis

Implementation and Usage

Common Use Cases

Caveats

Summary

- This talk is about inspecting streams of data for interesting patterns, even when that stream of data has been encoded.
- We focus on the Base64 encoding scheme, and discuss a tool that can be used when dealing with Base64.
- However, most portions of the algorithm are applicable to other position-dependent bitwise block encodings (and, potentially, self-synchronizing encodings).
- We also want to talk about how Erlang made development of the tool much easier.
 - This tool also provided an interesting “back door” to get Erlang accepted in DVLabs.

THE PROBLEM

In which we discover purpose of the whole thing...

SCREAM

Rob King

Introduction

The Problem

The Base64

Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression

Optimization

Performance

Analysis

Implementation and Usage

Common Use Cases

Caveats

Summary

- This talk is about inspecting streams of data for interesting patterns, even when that stream of data has been encoded.
- We focus on the Base64 encoding scheme, and discuss a tool that can be used when dealing with Base64.
- However, most portions of the algorithm are applicable to other position-dependent bitwise block encodings (and, potentially, self-synchronizing encodings).
- We also want to talk about how Erlang made development of the tool much easier.
 - This tool also provided an interesting “back door” to get Erlang accepted in DVLabs.

THE PROBLEM

In which we discover purpose of the whole thing...

SCREAM

Rob King

Introduction

The Problem

The Base64

Algorithm

Regular Expressions

The Algorithm

Ways of Solving the Problem

Encoding Operations

Performance

Expression

Optimization

Performance

Analysis

Implementation and Usage

Common Use Cases

Caveats

Summary

- This talk is about inspecting streams of data for interesting patterns, even when that stream of data has been encoded.
- We focus on the Base64 encoding scheme, and discuss a tool that can be used when dealing with Base64.
- However, most portions of the algorithm are applicable to other position-dependent bitwise block encodings (and, potentially, self-synchronizing encodings).
- We also want to talk about how Erlang made development of the tool much easier.
 - This tool also provided an interesting “back door” to get Erlang accepted in DVLabs.

THE PROBLEM

In which we discover purpose of the whole thing...

SCREAM

Rob King

Introduction

The Problem

The Base64

Algorithm

Regular Expressions

The Algorithm

Ways of Solving the Problem

Encoding Operations

Performance

Expression

Optimization

Performance

Analysis

Implementation and Usage

Common Use Cases

Caveats

Summary

- This talk is about inspecting streams of data for interesting patterns, even when that stream of data has been encoded.
- We focus on the Base64 encoding scheme, and discuss a tool that can be used when dealing with Base64.
- However, most portions of the algorithm are applicable to other position-dependent bitwise block encodings (and, potentially, self-synchronizing encodings).
- We also want to talk about how Erlang made development of the tool much easier.
 - This tool also provided an interesting “back door” to get Erlang accepted in DVLabs.

THE PROBLEM

In which we discover purpose of the whole thing...

SCREAM

Rob King

Introduction

The Problem

The Base64

Algorithm

Regular Expressions

The Algorithm

Ways of Solving the Problem

Encoding Operations

Performance

Expression

Optimization

Performance

Analysis

Implementation and Usage

Common Use Cases

Caveats

Summary

- This talk is about inspecting streams of data for interesting patterns, even when that stream of data has been encoded.
- We focus on the Base64 encoding scheme, and discuss a tool that can be used when dealing with Base64.
- However, most portions of the algorithm are applicable to other position-dependent bitwise block encodings (and, potentially, self-synchronizing encodings).
- We also want to talk about how Erlang made development of the tool much easier.
 - This tool also provided an interesting “back door” to get Erlang accepted in DVLabs.

CHALLENGES OF DATA STREAM INSPECTION

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
Optimization

Performance
Analysis

Implementation and Usage

Common Use Cases

Caveats

Summary

When looking for patterns in streams of data, several things must be kept in mind:

- There is no “luxury of time”.
- Context is limited.
- Resources are limited.

CHALLENGES OF DATA STREAM INSPECTION

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
Optimization
Performance
Analysis

Implementation and Usage

Common Use Cases
Caveats

Summary

When looking for patterns in streams of data, several things must be kept in mind:

- There is no “luxury of time”.
- Context is limited.
- Resources are limited.

CHALLENGES OF DATA STREAM INSPECTION

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
Optimization
Performance
Analysis

Implementation and Usage

Common Use Cases
Caveats

Summary

When looking for patterns in streams of data, several things must be kept in mind:

- There is no “luxury of time”.
- Context is limited.
- Resources are limited.

CHALLENGES OF DATA STREAM INSPECTION

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
Optimization
Performance
Analysis

Implementation and Usage

Common Use Cases
Caveats

Summary

When looking for patterns in streams of data, several things must be kept in mind:

- There is no “luxury of time”.
- Context is limited.
- Resources are limited.

SCREAM

Rob King

Introduction

The Problem

The Base64

Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression

Optimization

Performance

Analysis

Implementation
and Usage

Common Use Cases

Caveats

Summary

- Sometimes, the streams we're inspecting will be encoded.
- This means that we're going to have to be (more!) clever when looking for patterns in these streams.

SCREAM

Rob King

Introduction

The Problem

The Base64

Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression

Optimization

Performance

Analysis

Implementation
and Usage

Common Use Cases

Caveats

Summary

There are several general strategies for dealing with encoded streams.

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
Optimization

Performance
Analysis

Implementation and Usage

Common Use Cases

Caveats

Summary

- The easiest thing to do is simply pretend the stream is not encoded at all.
- Advantages:
 - We're already done.
- Disadvantages:
 - We're essentially admitting defeat.
 - Whatever we were looking for is not going to be found.
 - We're still burdening our analysis engine with lots of data with which we can do nothing.

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
Optimization

Performance
Analysis

Implementation and Usage

Common Use Cases

Caveats

Summary

- The easiest thing to do is simply pretend the stream is not encoded at all.
- Advantages:
 - We're already done.
- Disadvantages:
 - We're essentially admitting defeat.
 - Whatever we were looking for is not going to be found.
 - We're still burdening our analysis engine with lots of data with which we can do nothing.

SCREAM

Rob King

Introduction

The Problem

The Base64

Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
Optimization

Performance
Analysis

Implementation and Usage

Common Use Cases

Caveats

Summary

- The easiest thing to do is simply pretend the stream is not encoded at all.
- Advantages:
 - We're already done.
- Disadvantages:
 - We're essentially admitting defeat.
 - Whatever we were looking for is not going to be found.
 - We're still burdening our analysis engine with lots of data with which we can do nothing.

SCREAM

Rob King

Introduction

The Problem

The Base64

Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression

Optimization

Performance

Analysis

Implementation and Usage

Common Use Cases

Caveats

Summary

- Rather than pretending that the data is not encoded, we could go one step further, and detect that the data is encoded.
- Once we've detected that the data is encoded, we can simply drop the stream.
- Advantages:
 - Stops burdening with inspection engine with data we know we can't inspect.
- Disadvantages:
 - If we “fail open” and allow encoded data to pass without inspection, we just gave anyone who wants to bypass our inspection a “get out of jail free” card.
 - If we “fail closed” and block encoded data, we just blocked all legitimate uses of that encoding scheme.

SCREAM

Rob King

Introduction

The Problem

The Base64

Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression

Optimization

Performance

Analysis

Implementation and Usage

Common Use Cases

Caveats

Summary

- Rather than pretending that the data is not encoded, we could go one step further, and detect that the data is encoded.
- Once we've detected that the data is encoded, we can simply drop the stream.
- Advantages:
 - Stops burdening with inspection engine with data we know we can't inspect.
- Disadvantages:
 - If we “fail open” and allow encoded data to pass without inspection, we just gave anyone who wants to bypass our inspection a “get out of jail free” card.
 - If we “fail closed” and block encoded data, we just blocked all legitimate uses of that encoding scheme.

SCREAM

Rob King

Introduction

The Problem

The Base64

Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression

Optimization

Performance

Analysis

Implementation and Usage

Common Use Cases

Caveats

Summary

- Rather than pretending that the data is not encoded, we could go one step further, and detect that the data is encoded.
- Once we've detected that the data is encoded, we can simply drop the stream.
- Advantages:
 - Stops burdening with inspection engine with data we know we can't inspect.
- Disadvantages:
 - If we "fail open" and allow encoded data to pass without inspection, we just gave anyone who wants to bypass our inspection a "get out of jail free" card.
 - If we "fail closed" and block encoded data, we just blocked all legitimate uses of that encoding scheme.

SCREAM

Rob King

Introduction

The Problem

The Base64

Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression

Optimization

Performance

Analysis

Implementation

and Usage

Common Use Cases

Caveats

Summary

- Rather than pretending that the data is not encoded, we could go one step further, and detect that the data is encoded.
- Once we've detected that the data is encoded, we can simply drop the stream.
- Advantages:
 - Stops burdening with inspection engine with data we know we can't inspect.
- Disadvantages:
 - If we “fail open” and allow encoded data to pass without inspection, we just gave anyone who wants to bypass our inspection a “get out of jail free” card.
 - If we “fail closed” and block encoded data, we just blocked all legitimate uses of that encoding scheme.

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
Optimization

Performance
Analysis

Implementation and Usage

Common Use Cases

Caveats

Summary

- We could buffer the entirety of an encoded stream.
- Once we've buffered the whole stream, we can decode it, inspect it, reencode it, and send it on its way.
- Advantages:
 - We get the complete power of our inspection engine.
- Disadvantages:
 - Latency becomes unbounded.
 - Resource usage becomes unbounded.
 - We have to modify the engine for every encoding we need to inspect.
- These advantages and disadvantages apply roughly equally to any streaming decoder.

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
Optimization

Performance
Analysis

Implementation and Usage

Common Use Cases

Caveats

Summary

- We could buffer the entirety of an encoded stream.
- Once we've buffered the whole stream, we can decode it, inspect it, reencode it, and send it on its way.
- Advantages:
 - We get the complete power of our inspection engine.
- Disadvantages:
 - Latency becomes unbounded.
 - Resource usage becomes unbounded.
 - We have to modify the engine for every encoding we need to inspect.
- These advantages and disadvantages apply roughly equally to any streaming decoder.

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
OptimizationPerformance
AnalysisImplementation
and Usage

Common Use Cases

Caveats

Summary

- We could buffer the entirety of an encoded stream.
- Once we've buffered the whole stream, we can decode it, inspect it, reencode it, and send it on its way.
- Advantages:
 - We get the complete power of our inspection engine.
- Disadvantages:
 - Latency becomes unbounded.
 - Resource usage becomes unbounded.
 - We have to modify the engine for every encoding we need to inspect.
- These advantages and disadvantages apply roughly equally to any streaming decoder.

SCREAM

Rob King

Introduction

The Problem

The Base64

Algorithm

Regular Expressions

The Algorithm

Ways of Solving the

Problem

Encoding Operations

Performance

Expression

Optimization

Performance

Analysis

Implementation and Usage

Common Use Cases

Caveats

Summary

- We could buffer the entirety of an encoded stream.
- Once we've buffered the whole stream, we can decode it, inspect it, reencode it, and send it on its way.
- Advantages:
 - We get the complete power of our inspection engine.
- Disadvantages:
 - Latency becomes unbounded.
 - Resource usage becomes unbounded.
 - We have to modify the engine for every encoding we need to inspect.
- These advantages and disadvantages apply roughly equally to any streaming decoder.

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
OptimizationPerformance
AnalysisImplementation
and Usage

Common Use Cases

Caveats

Summary

- We could buffer the entirety of an encoded stream.
- Once we've buffered the whole stream, we can decode it, inspect it, reencode it, and send it on its way.
- Advantages:
 - We get the complete power of our inspection engine.
- Disadvantages:
 - Latency becomes unbounded.
 - Resource usage becomes unbounded.
 - We have to modify the engine for every encoding we need to inspect.
- These advantages and disadvantages apply roughly equally to any streaming decoder.

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
Optimization

Performance
Analysis

Implementation and Usage

Common Use Cases
Caveats

Summary

- Instead of decoding the data, we could encode the pattern.
- Advantages:
 - We get (most?) of the power of our inspection engine.
 - There is no performance penalty for decoding.
 - Resource usage and latency are bounded.
 - We need not buffer or store context.
- Disadvantages:
 - A different transform must be written for each scheme.
 - Might not be possible for some schemes or patterns.
 - False positives may become more common.

SCREAM

Rob King

Introduction

The Problem

The Base64

Algorithm

Regular Expressions

The Algorithm

Ways of Solving the Problem

Encoding Operations

Performance

Expression

Optimization

Performance

Analysis

Implementation and Usage

Common Use Cases

Caveats

Summary

- Instead of decoding the data, we could encode the pattern.
- Advantages:
 - We get (most?) of the power of our inspection engine.
 - There is no performance penalty for decoding.
 - Resource usage and latency are bounded.
 - We need not buffer or store context.
- Disadvantages:
 - A different transform must be written for each scheme.
 - Might not be possible for some schemes or patterns.
 - False positives may become more common.

SCREAM

Rob King

Introduction

The Problem

The Base64

Algorithm

Regular Expressions

The Algorithm

Ways of Solving the

Problem

Encoding Operations

Performance

Expression

Optimization

Performance

Analysis

Implementation
and Usage

Common Use Cases

Caveats

Summary

- Instead of decoding the data, we could encode the pattern.
- Advantages:
 - We get (most?) of the power of our inspection engine.
 - There is no performance penalty for decoding.
 - Resource usage and latency are bounded.
 - We need not buffer or store context.
- Disadvantages:
 - A different transform must be written for each scheme.
 - Might not be possible for some schemes or patterns.
 - False positives may become more common.

BASE64

In which we discover the usefulness of large radices or “Maybe the Sumerians were right after all...”

SCREAM

Rob King

Introduction

The Problem

The Base64 Algorithm

Regular Expressions

The Algorithm

Ways of Solving the Problem

Encoding Operations

Performance

Expression

Optimization

Performance

Analysis

Implementation and Usage

Common Use Cases

Caveats

Summary

- Base64 is an Internet standard for encoding arbitrary (usually binary) data using only printable characters common in many character sets.
- Multiple minor variants, but the most common is defined in RFC4648.
- Base64 is used in numerous situations, essentially whenever binary data needs to be encoded in a printable form.

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
OptimizationPerformance
AnalysisImplementation
and Usage

Common Use Cases

Caveats

Summary

- Base64 expects input as a series of eight-bit octets.
- Every three octets are grouped together into a collection of 24-bits.
- These 24-bits are then split into four six-bit sextets.
- Each sextet is used as a big-endian index into the zero-based array that is the Base64 alphabet.

f								o								o															
0	1	1	0	0	1	1	0	0	1	1	0	1	1	1	1	0	1	1	0	1	1	1	1								
Z								m								9								v							

- If the input is not evenly divisible by three, it is right-padded by one or two octets of zeroes, and one or two "=" symbols are appended to the output.

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
OptimizationPerformance
AnalysisImplementation
and Usage

Common Use Cases

Caveats

Summary

- Base64 expects input as a series of eight-bit octets.
- Every three octets are grouped together into a collection of 24-bits.
 - These 24-bits are then split into four six-bit sextets.
 - Each sextet is used as a big-endian index into the zero-based array that is the Base64 alphabet.

f								o								o															
0	1	1	0	0	1	1	0	0	1	1	0	1	1	1	1	0	1	1	0	1	1	1	1								
Z								m								9								v							

- If the input is not evenly divisible by three, it is right-padded by one or two octets of zeroes, and one or two "=" symbols are appended to the output.

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
OptimizationPerformance
AnalysisImplementation
and Usage

Common Use Cases

Caveats

Summary

- Base64 expects input as a series of eight-bit octets.
- Every three octets are grouped together into a collection of 24-bits.
- These 24-bits are then split into four six-bit sextets.
- Each sextet is used as a big-endian index into the zero-based array that is the Base64 alphabet.

f								o								o															
0	1	1	0	0	1	1	0	0	1	1	0	1	1	1	1	0	1	1	0	1	1	1	1								
Z								m								9								v							

- If the input is not evenly divisible by three, it is right-padded by one or two octets of zeroes, and one or two "=" symbols are appended to the output.

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
OptimizationPerformance
AnalysisImplementation
and Usage

Common Use Cases

Caveats

Summary

- Base64 expects input as a series of eight-bit octets.
- Every three octets are grouped together into a collection of 24-bits.
- These 24-bits are then split into four six-bit sextets.
- Each sextet is used as a big-endian index into the zero-based array that is the Base64 alphabet.

f								o								o															
0	1	1	0	0	1	1	0	0	1	1	0	1	1	1	1	0	1	1	0	1	1	1	1								
Z								m								9								v							

- If the input is not evenly divisible by three, it is right-padded by one or two octets of zeroes, and one or two "=" symbols are appended to the output.

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
OptimizationPerformance
AnalysisImplementation
and Usage

Common Use Cases

Caveats

Summary

- Base64 expects input as a series of eight-bit octets.
- Every three octets are grouped together into a collection of 24-bits.
- These 24-bits are then split into four six-bit sextets.
- Each sextet is used as a big-endian index into the zero-based array that is the Base64 alphabet.

f								o								o															
0	1	1	0	0	1	1	0	0	1	1	0	1	1	1	1	0	1	1	0	1	1	1	1								
Z								m								9								v							

- If the input is not evenly divisible by three, it is right-padded by one or two octets of zeroes, and one or two "=" symbols are appended to the output.

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
OptimizationPerformance
AnalysisImplementation
and Usage

Common Use Cases

Caveats

Summary

- Base64 expects input as a series of eight-bit octets.
- Every three octets are grouped together into a collection of 24-bits.
- These 24-bits are then split into four six-bit sextets.
- Each sextet is used as a big-endian index into the zero-based array that is the Base64 alphabet.

f								o							o													
0	1	1	0	0	1	1	0	0	1	1	0	1	1	1	1	0	1	1	0	1	1	1	1					
Z								m							9							v						

- If the input is not evenly divisible by three, it is right-padded by one or two octets of zeroes, and one or two "=" symbols are appended to the output.

BASE64

The Base64 Alphabet

SCREAM

Rob King

Introduction

The Problem

**The Base64
Algorithm**

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
OptimizationPerformance
Analysis

Implementation and Usage

Common Use Cases

Caveats

Summary

	0	1	2	3	4	5	6	7	8	9	A	B	C	D	E	F
0	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P
1	Q	R	S	T	U	V	W	X	Y	Z	a	b	c	d	e	f
2	g	h	i	j	k	l	m	n	o	p	q	r	s	t	u	v
3	w	x	y	z	0	1	2	3	4	5	6	7	8	9	+	/
The “=” sign is used for padding.																

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
OptimizationPerformance
AnalysisImplementation
and Usage

Common Use Cases

Caveats

Summary

- Base64 is position dependent.
- In general, any given string will be encoded in one of three different ways, depending on its offset into the input.
- This is what makes the encoding of patterns so hard.

Encoding of "foo" at Three Different Offsets	
Offset 0	Zm9v
Offset 1	[159BFJNRVZdhlptx]mb2[+/8-9]
Offset 2	[2GWm]Zvb[+/0-9w-z]

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
Optimization

Performance
Analysis

Implementation and Usage

Common Use Cases

Caveats

Summary

- Our choice of pattern language controls the overall complexity of patterns for which we can search.
- If we were just looking for static strings, we wouldn't need a new tool - just use grep.
- Regular expressions provide a good balance of ease of implementation, expressive power, and common availability.

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
Optimization

Performance
Analysis

Implementation and Usage

Common Use Cases

Caveats

Summary

- Formalized by Stephen Kleene in 1956.
- Ken Thompson incorporated them as a useful pattern matching tool into his version of the QED editor for MIT's CTSS timesharing system.
- This later influenced Ken Thompson's implementation of `ed` for UNIX.
- From UNIX, regular expressions spread around the world.

SUPPORTED REGULAR EXPRESSION OPERATIONS

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
Optimization

Performance
Analysis

Implementation and Usage

Common Use Cases

Caveats

Summary

- Character matches (e.g. “a”)
- Concatenation (e.g. “ab”)
- Alternation (e.g. “(ab|cd)”)
- Character classes and inverse classes (e.g. “[0 – 9]”)
- Kleene closures (e.g. “A*C”)

UNSUPPORTED REGULAR EXPRESSION OPERATIONS

SCREAM

Rob King

Introduction

The Problem

The Base64

Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression

Optimization

Performance

Analysis

Implementation and Usage

Common Use Cases

Caveats

Summary

- Backreferences and captures
- Variable-length repetition
- Left and right anchors
- Just about everything else

SCREAM

Rob King

Introduction

The Problem

The Base64

Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression

Optimization

Performance

Analysis

Implementation
and Usage

Common Use Cases

Caveats

Summary

- Come up with a way to transform an arbitrary regular expression such that it will match its input when that input has been encoded using the Base64 algorithm.
- Do this transformation in such a way that the regular expression does not grow too large.
- Do this transformation in such a way that not too much information is lost.
- Do this transformation in such a way that the expression will match regardless of the pattern's offset from the beginning of input.

SCREAM

Rob King

Introduction

The Problem

The Base64

Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression

Optimization

Performance

Analysis

Implementation
and Usage

Common Use Cases

Caveats

Summary

- Come up with a way to transform an arbitrary regular expression such that it will match its input when that input has been encoded using the Base64 algorithm.
- Do this transformation in such a way that the regular expression does not grow too large.
- Do this transformation in such a way that not too much information is lost.
- Do this transformation in such a way that the expression will match regardless of the pattern's offset from the beginning of input.

SCREAM

Rob King

Introduction

The Problem

The Base64

Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression

Optimization

Performance

Analysis

Implementation
and Usage

Common Use Cases

Caveats

Summary

- Come up with a way to transform an arbitrary regular expression such that it will match its input when that input has been encoded using the Base64 algorithm.
- Do this transformation in such a way that the regular expression does not grow too large.
- Do this transformation in such a way that not too much information is lost.
- Do this transformation in such a way that the expression will match regardless of the pattern's offset from the beginning of input.

SCREAM

Rob King

Introduction

The Problem

The Base64

Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression

Optimization

Performance

Analysis

Implementation
and Usage

Common Use Cases

Caveats

Summary

- Come up with a way to transform an arbitrary regular expression such that it will match its input when that input has been encoded using the Base64 algorithm.
- Do this transformation in such a way that the regular expression does not grow too large.
- Do this transformation in such a way that not too much information is lost.
- Do this transformation in such a way that the expression will match regardless of the pattern's offset from the beginning of input.

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
Optimization

Performance
Analysis

Implementation and Usage

Common Use Cases

Caveats

Summary

- Convert the regular expression into a nondeterministic finite state automaton using Thompson's algorithm, then convert the NFA to a deterministic finite state automaton using the powerset construction algorithm, then transform the DFA into a directed acyclic graph, then perform graph reductions until the graph is in a minimal form, transform the minimal form using graph transformations, then reduce again, then serialize the graph as a regular expression.
- This is hard, and I'm lazy.

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
OptimizationPerformance
AnalysisImplementation
and Usage

Common Use Cases

Caveats

Summary

- Convert the regular expression into a nondeterministic finite state automaton using Thompson's algorithm, then convert the NFA to a deterministic finite state automaton using the powerset construction algorithm, then transform the DFA into a directed acyclic graph, then perform graph reductions until the graph is in a minimal form, transform the minimal form using graph transformations, then reduce again, then serialize the graph as a regular expression.
- This is hard, and I'm lazy.

SCREAM

Rob King

Introduction

The Problem

The Base64

Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression

Optimization

Performance

Analysis

Implementation
and Usage

Common Use Cases

Caveats

Summary

- Convert the regular expression into a nondeterministic finite state automaton using Thompson's algorithm, then convert the NFA to a deterministic finite state automaton using the powerset construction algorithm, then transform the DFA into a directed acyclic graph, then perform graph reductions until the graph is in a minimal form, transform the minimal form using graph transformations, then reduce again, then serialize the graph as a regular expression.
- This is hard, and I'm lazy.

SCREAM

Rob King

Introduction

The Problem

The Base64

Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression

Optimization

Performance

Analysis

Implementation and Usage

Common Use Cases

Caveats

Summary

- Enumerate all possible matching strings for a regular expression.
- Encode these strings and concatenate them inside an alternating regular expression.
- This would definitely work, with one problem...
- Regular expressions with Kleene closures can match an infinite number of strings.
- Enumerating an infinite number of strings can take a really long time.
- Even the trivial regular expression “...” would match over four billion strings.

SCREAM

Rob King

Introduction

The Problem

The Base64

Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression

Optimization

Performance

Analysis

Implementation
and Usage

Common Use Cases

Caveats

Summary

- Enumerate all possible matching strings for a regular expression.
- Encode these strings and concatenate them inside an alternating regular expression.
- This would definitely work, with one problem...
- Regular expressions with Kleene closures can match an infinite number of strings.
- Enumerating an infinite number of strings can take a really long time.
- Even the trivial regular expression “...” would match over four billion strings.

SCREAM

Rob King

Introduction

The Problem

The Base64

Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression

Optimization

Performance

Analysis

Implementation
and Usage

Common Use Cases

Caveats

Summary

- Enumerate all possible matching strings for a regular expression.
- Encode these strings and concatenate them inside an alternating regular expression.
- This would definitely work, with one problem...
- Regular expressions with Kleene closures can match an infinite number of strings.
- Enumerating an infinite number of strings can take a really long time.
- Even the trivial regular expression “...” would match over four billion strings.

SCREAM

Rob King

Introduction

The Problem

The Base64

Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression

Optimization

Performance

Analysis

Implementation
and Usage

Common Use Cases

Caveats

Summary

- Enumerate all possible matching strings for a regular expression.
- Encode these strings and concatenate them inside an alternating regular expression.
- This would definitely work, with one problem...
- Regular expressions with Kleene closures can match an infinite number of strings.
- Enumerating an infinite number of strings can take a really long time.
- Even the trivial regular expression “...” would match over four billion strings.

SCREAM

Rob King

Introduction

The Problem

The Base64

Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression

Optimization

Performance

Analysis

Implementation
and Usage

Common Use Cases

Caveats

Summary

- Enumerate all possible matching strings for a regular expression.
- Encode these strings and concatenate them inside an alternating regular expression.
- This would definitely work, with one problem...
- Regular expressions with Kleene closures can match an infinite number of strings.
- Enumerating an infinite number of strings can take a really long time.
- Even the trivial regular expression “...” would match over four billion strings.

SCREAM

Rob King

Introduction

The Problem

The Base64

Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression

Optimization

Performance

Analysis

Implementation
and Usage

Common Use Cases

Caveats

Summary

- Enumerate all possible matching strings for a regular expression.
- Encode these strings and concatenate them inside an alternating regular expression.
- This would definitely work, with one problem...
- Regular expressions with Kleene closures can match an infinite number of strings.
- Enumerating an infinite number of strings can take a really long time.
- Even the trivial regular expression “...” would match over four billion strings.

SCREAM

Rob King

Introduction

The Problem

The Base64

Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression

Optimization

Performance

Analysis

Implementation
and Usage

Common Use Cases

Caveats

Summary

- Enumerate all possible matching strings for a regular expression.
- Encode these strings and concatenate them inside an alternating regular expression.
- This would definitely work, with one problem...
- Regular expressions with Kleene closures can match an infinite number of strings.
- Enumerating an infinite number of strings can take a really long time.
- Even the trivial regular expression “...” would match over four billion strings.

WAYS OF SOLVING THE PROBLEM

The Wrong Way, But Faster!

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
Optimization

Performance
Analysis

Implementation and Usage

Common Use Cases

Caveats

Summary

- Enumerating all possible matching strings would, in theory, give us the correct answer.
- Instead of enumerating all possible strings and encoding them, what if we encoded the operations?
- In other words, what if we enumerated only what could match a given operation at a given point in time?

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
Optimization

Performance
Analysis

Implementation and Usage

Common Use Cases

Caveats

Summary

- Enumerating all possible matching strings would, in theory, give us the correct answer.
- Instead of enumerating all possible strings and encoding them, what if we encoded the operations?
- In other words, what if we enumerated only what could match a given operation at a given point in time?

WAYS OF SOLVING THE PROBLEM

The Wrong Way, But Faster!

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
Optimization

Performance
Analysis

Implementation and Usage

Common Use Cases

Caveats

Summary

- Enumerating all possible matching strings would, in theory, give us the correct answer.
- Instead of enumerating all possible strings and encoding them, what if we encoded the operations?
- In other words, what if we enumerated only what could match a given operation at a given point in time?

WAYS OF SOLVING THE PROBLEM

The Wrong Way, But Faster!

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
Optimization

Performance
Analysis

Implementation and Usage

Common Use Cases

Caveats

Summary

- Enumerating all possible matching strings would, in theory, give us the correct answer.
- Instead of enumerating all possible strings and encoding them, what if we encoded the operations?
- In other words, what if we enumerated only what could match a given operation at a given point in time?

SCREAM

Rob King

Introduction

The Problem

The Base64

Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression

Optimization

Performance

Analysis

Implementation
and Usage

Common Use Cases

Caveats

Summary

- Create a list of regular expressions that match a fixed-length string.
- Enumerate the bitstrings matched by each of these expressions.
- Break each of these bitstrings into six-bit units.
- Encode each of these units.
- Treat each of these encoded strings as a branch in an n -way alternation in a regular expression.
- Optimize the expression.
- There is special handling for the two “meta” operations: alternations and closures.

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
OptimizationPerformance
AnalysisImplementation
and Usage

Common Use Cases

Caveats

Summary

- Create a list of regular expressions that match a fixed-length string.
- Enumerate the bitstrings matched by each of these expressions.
- Break each of these bitstrings into six-bit units.
- Encode each of these units.
- Treat each of these encoded strings as a branch in an n -way alternation in a regular expression.
- Optimize the expression.
- There is special handling for the two “meta” operations: alternations and closures.

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
OptimizationPerformance
AnalysisImplementation
and Usage

Common Use Cases

Caveats

Summary

- Create a list of regular expressions that match a fixed-length string.
- Enumerate the bitstrings matched by each of these expressions.
 - Break each of these bitstrings into six-bit units.
 - Encode each of these units.
 - Treat each of these encoded strings as a branch in an n -way alternation in a regular expression.
 - Optimize the expression.
 - There is special handling for the two “meta” operations: alternations and closures.

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
OptimizationPerformance
AnalysisImplementation
and Usage

Common Use Cases

Caveats

Summary

- Create a list of regular expressions that match a fixed-length string.
- Enumerate the bitstrings matched by each of these expressions.
- Break each of these bitstrings into six-bit units.
- Encode each of these units.
- Treat each of these encoded strings as a branch in an n -way alternation in a regular expression.
- Optimize the expression.
- There is special handling for the two “meta” operations: alternations and closures.

SCREAM

Rob King

Introduction

The Problem

The Base64

Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression

Optimization

Performance

Analysis

Implementation
and Usage

Common Use Cases

Caveats

Summary

- Create a list of regular expressions that match a fixed-length string.
- Enumerate the bitstrings matched by each of these expressions.
- Break each of these bitstrings into six-bit units.
- Encode each of these units.
- Treat each of these encoded strings as a branch in an n -way alternation in a regular expression.
- Optimize the expression.
- There is special handling for the two “meta” operations: alternations and closures.

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
OptimizationPerformance
AnalysisImplementation
and Usage

Common Use Cases

Caveats

Summary

- Create a list of regular expressions that match a fixed-length string.
- Enumerate the bitstrings matched by each of these expressions.
- Break each of these bitstrings into six-bit units.
- Encode each of these units.
- Treat each of these encoded strings as a branch in an n -way alternation in a regular expression.
- Optimize the expression.
- There is special handling for the two “meta” operations: alternations and closures.

SCREAM

Rob King

Introduction

The Problem

The Base64

Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression

Optimization

Performance

Analysis

Implementation
and Usage

Common Use Cases

Caveats

Summary

- Create a list of regular expressions that match a fixed-length string.
- Enumerate the bitstrings matched by each of these expressions.
- Break each of these bitstrings into six-bit units.
- Encode each of these units.
- Treat each of these encoded strings as a branch in an n -way alternation in a regular expression.
- Optimize the expression.
- There is special handling for the two “meta” operations: alternations and closures.

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
OptimizationPerformance
AnalysisImplementation
and Usage

Common Use Cases

Caveats

Summary

- Create a list of regular expressions that match a fixed-length string.
- Enumerate the bitstrings matched by each of these expressions.
- Break each of these bitstrings into six-bit units.
- Encode each of these units.
- Treat each of these encoded strings as a branch in an n -way alternation in a regular expression.
- Optimize the expression.
- There is special handling for the two “meta” operations: alternations and closures.

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
Optimization

Performance
Analysis

Implementation and Usage

Common Use Cases

Caveats

Summary

- Non-fixed-length operations leave ambiguity as to what bits to place in a given six-bit unit.
- The goal of this phase of the algorithm is to get a collection of regular expressions that match fixed-length strings.
- First, let's go over how to encode the atomic operations.

SCREAM

Rob King

Introduction

The Problem

The Base64

Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression

Optimization

Performance

Analysis


Implementation
and Usage

Common Use Cases

Caveats

Summary

This is a fairly obvious - a single character is encoded as the bitstring for that character.



SCREAM

Rob King

Introduction

The Problem

The Base64

Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression

Optimization

Performance

Analysis

Implementation
and Usage

Common Use Cases

Caveats

Summary

Character classes are stored as a list of all characters included in the class. This is not the most efficient possible choice in terms of memory usage, but its implementation is simpler and it is often faster.

[ABC]							
0	1	0	0	0	0	0	1
0	1	0	0	0	0	1	0
0	1	0	0	0	0	1	1

Inverted character classes simply store the complement of the list of indicated characters.

SCREAM

Rob King

Introduction

- The Problem
- The Base64 Algorithm
- Regular Expressions

The Algorithm

- Ways of Solving the Problem
- Encoding Operations

Performance

- Expression Optimization
- Performance Analysis

Implementation and Usage

- Common Use Cases
- Caveats

Summary

Wildcards are simply encoded as a character class with 256 entries. In the underlying implementation they are encoded as a special atom to save space, but from the point of view of the algorithm there is no difference.

SCREAM

Rob King

Introduction

The Problem

The Base64

Algorithm

Regular Expressions

The Algorithm

Ways of Solving the

Problem

Encoding Operations

Performance

Expression

Optimization

Performance

Analysis


Implementation
and Usage

Common Use Cases

Caveats

Summary

- Single characters, character classes, and wildcards are treated as the building blocks of lists.
- Each fixed-length expression is simply a list of these operations.
- Therefore, concatenation is modelled implicitly by the ordering of these operations in the list.



SCREAM

Rob King

Introduction

The Problem

The Base64

Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression

Optimization

Performance

Analysis

Implementation
and Usage

Common Use Cases

Caveats

Summary

- It was stated earlier that the algorithm deals only with regular expressions that match only fixed-length strings.
- Alternations can easily violate this: if one branch of an alternation has more character matches than the other, the expression overall is not fixed length.

SCREAM

Rob King

Introduction

The Problem

The Base64

Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression

Optimization

Performance

Analysis

Implementation
and Usage

Common Use Cases

Caveats

Summary

- It was stated earlier that the algorithm deals only with regular expressions that match only fixed-length strings.
- Alternations can easily violate this: if one branch of an alternation has more character matches than the other, the expression overall is not fixed length.

SCREAM

Rob King

Introduction

The Problem

The Base64

Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression

Optimization

Performance

Analysis

Implementation
and Usage

Common Use Cases

Caveats

Summary

- It was stated earlier that the algorithm deals only with regular expressions that match only fixed-length strings.
- Alternations can easily violate this: if one branch of an alternation has more character matches than the other, the expression overall is not fixed length.

SCREAM

Rob King

Introduction

The Problem

The Base64

Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression

Optimization

Performance

Analysis


Implementation
and Usage

Common Use Cases

Caveats

Summary

- Alternations turn our simple lists of operations into directed acyclic graphs.
- For example, below is the representation of the regular expression “A(B|(C|D)E)F”.



SCREAM

Rob King

Introduction

The Problem

The Base64

Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression

Optimization

Performance

Analysis

Implementation
and Usage


Common Use Cases

Caveats

Summary

- We can easily create a list of fixed-length expressions by enumerating all topological sorts of the graph.
- For each distinct topological sort, we can create a list of fixed-length operations.
- This reduces an alternation to a fixed-length structure, and can be performed recursively, resulting in a list of expressions, each of a fixed-length.

A(BI(CID)E)F



SCREAM

Rob King

Introduction

The Problem

The Base64

Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression

Optimization

Performance

Analysis


Implementation
and Usage

Common Use Cases

Caveats

Summary

- By adding an explicit “end of expression” pseudo-operation to the end of every expression, we can guarantee that all expressions have a shared tail.
- We can also add some number of explicit “start of expression” pseudo-operations from every entry point into the expression.
- These operations does not actually match anything; they exist solely for ease of implementation.



SCREAM

Rob King

Introduction

The Problem

The Base64

Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression

Optimization

Performance

Analysis

Implementation
and Usage

Common Use Cases

Caveats

Summary

- By finding a path from every START node to the END node, we can produce a list of fixed-length expressions.
- Note that this is essentially equivalent of the list of topological sorts, but provides the storage advantage of a shared tail.
- Implementation is also considerably easier and performs much better, since we don't need to implement a general topological sort; rather we can simply maintain a list of all start nodes and follow the only emanating edge from every node until we reach the end.

SCREAM

Rob King

Introduction

The Problem

The Base64

Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression

Optimization

Performance

Analysis


Implementation
and Usage

Common Use Cases

Caveats

Summary

- We can apply this recursively, to build up a list of fixed-length expressions from a regular expression, even with nested alternations.
- Each of these fixed length expressions can be encoded easily using the methods for the fixed-length operations above.



ENCODING EXPRESSIONS WITHOUT CLOSURES

Create a List of Fixed Length Bit Expressions

SCREAM

Rob King

Introduction

The Problem

The Base64

Algorithm

Regular Expressions

The Algorithm

Ways of Solving the Problem

Encoding Operations

Performance

Expression

Optimization

Performance

Analysis


Implementation and Usage

Common Use Cases

Caveats

Summary

- Using the algorithm above, we can create a list of fixed-length expressions consisting of eight bit units (for single characters) or lists of eight bit units (for character classes and wildcards).



ENCODING EXPRESSIONS WITHOUT CLOSURES

The Initial and Final Bits

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
Optimization

Performance
Analysis

Implementation and Usage

Common Use Cases

Caveats

Summary

- Since the expression is going to be examining data in a streaming context, left and right anchors don't make sense in most situations.
- Therefore, if any bits are needed to complete a six-bit unit at the beginning or the end of the expression, they are treated as if a wildcard were present immediately before or after the expression.

ENCODING EXPRESSIONS WITHOUT CLOSURES

The Initial and Final Bits

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
Optimization

Performance
Analysis

Implementation and Usage

Common Use Cases

Caveats

Summary

- Since the expression is going to be examining data in a streaming context, left and right anchors don't make sense in most situations.
- Therefore, if any bits are needed to complete a six-bit unit at the beginning or the end of the expression, they are treated as if a wildcard were present immediately before or after the expression.

ENCODING EXPRESSIONS WITHOUT CLOSURES

The Initial and Final Bits

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
Optimization

Performance
Analysis

Implementation and Usage

Common Use Cases

Caveats

Summary

- Since the expression is going to be examining data in a streaming context, left and right anchors don't make sense in most situations.
- Therefore, if any bits are needed to complete a six-bit unit at the beginning or the end of the expression, they are treated as if a wildcard were present immediately before or after the expression.

ENCODING EXPRESSIONS WITHOUT CLOSURES

Encode the Expressions

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
Optimization

Performance
Analysis


Implementation
and Usage

Common Use Cases

Caveats

Summary

- Each six bit unit is then replaced with its equivalent symbol in the Base64 alphabet.
- For lists of six bit units (from character classes), a list of Base64 symbols is produced.



ENCODING EXPRESSIONS WITHOUT CLOSURES

The Final Expression

SCREAM

Rob King

Introduction

The Problem

The Base64 Algorithm

Regular Expressions

The Algorithm

Ways of Solving the Problem

Encoding Operations

Performance

Expression Optimization

Performance Analysis

Implementation and Usage

Common Use Cases

Caveats

Summary

- The list of fixed-length encoded expressions is joined into an n -way alternation, which is treated as a regular expression.
- This expression is first run through the optimizer, which recursively refactors common prefixes and suffixes.
- The expression is then reprocessed two more times, to account for varying offsets from the beginning of input.

ENCODING EXPRESSIONS WITH CLOSURES

Violating the Fixed Length and Acyclic Rules

SCREAM

Rob King

Introduction

The Problem

The Base64

Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression

Optimization

Performance

Analysis


Implementation
and Usage

Common Use Cases

Caveats

Summary

- Expressions with Kleene closures violate the fixed-length rule, since subexpressions can appear from zero to infinitely many times.
- Expressions with Kleene closures violate the acyclic rule, since repeating expressions would become cycles in the graph representation.



ENCODING EXPRESSIONS WITH CLOSURES

The Easy Case

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
Optimization

Performance
Analysis

Implementation and Usage

Common Use Cases

Caveats

Summary

- The easiest case for encoding closures is the case where they don't appear at all.
- Whenever a closure is present in an expression, it acts as a virtual alternation, with one branch having at least one instance of the closure, and the other branch having no instances.

ENCODING EXPRESSIONS WITH CLOSURES

Marking Closures

SCREAM

Rob King

Introduction

The Problem

The Base64

Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression

Optimization

Performance

Analysis

Implementation and Usage

Common Use Cases

Caveats

Summary

- Closures are encoded three times, to account for varying offsets from the beginning of input.
- These three encoded expressions are placed in a special data structure that marks them as the components of a closure.

ENCODING EXPRESSIONS WITH CLOSURES

A List of Expressions

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
Optimization

Performance
Analysis

Implementation and Usage

Common Use Cases

Caveats

Summary

- Expressions with closures are treated as a list of expressions.
- The expressions are split such that each specially-encoded closure is a separate expression.
- Lists of expressions are encoded such that the implicit wildcard at the beginning and ending of each expression is replaced with the appropriate values from the next or previous expression in the list.

SCREAM

Rob King

Introduction

The Problem

The Base64

Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
Optimization

Performance
Analysis

Implementation and Usage

Common Use Cases

Caveats

Summary

- The expression is treated as a list of constant-length expressions, from each possible starting expression to the logical end.
- All of these expressions will differ only around places where alternations were present.
- Therefore, everything leading up to the alternation, and everything after the alternation, will be identical for various expressions.
- We can therefore optimize the expression by refactoring all common prefixes and suffixes.

SCREAM

Rob King

Introduction

The Problem

The Base64

Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
Optimization

Performance
Analysis

Implementation and Usage

Common Use Cases

Caveats

Summary

- The expression is treated as a list of constant-length expressions, from each possible starting expression to the logical end.
- All of these expressions will differ only around places where alternations were present.
- Therefore, everything leading up to the alternation, and everything after the alternation, will be identical for various expressions.
- We can therefore optimize the expression by refactoring all common prefixes and suffixes.

SCREAM

Rob King

Introduction

The Problem

The Base64

Algorithm

Regular Expressions

The Algorithm

Ways of Solving the

Problem

Encoding Operations

Performance

Expression

Optimization

Performance

Analysis

Implementation and Usage

Common Use Cases

Caveats

Summary

- The expression is treated as a list of constant-length expressions, from each possible starting expression to the logical end.
- All of these expressions will differ only around places where alternations were present.
- Therefore, everything leading up to the alternation, and everything after the alternation, will be identical for various expressions.
- We can therefore optimize the expression by refactoring all common prefixes and suffixes.

SCREAM

Rob King

Introduction

The Problem

The Base64

Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
OptimizationPerformance
AnalysisImplementation
and Usage

Common Use Cases

Caveats

Summary

- The expression is treated as a list of constant-length expressions, from each possible starting expression to the logical end.
- All of these expressions will differ only around places where alternations were present.
- Therefore, everything leading up to the alternation, and everything after the alternation, will be identical for various expressions.
- We can therefore optimize the expression by refactoring all common prefixes and suffixes.

SCREAM

Rob King

Introduction

The Problem

The Base64

Algorithm

Regular Expressions

The Algorithm

Ways of Solving the

Problem

Encoding Operations

Performance

Expression

Optimization

Performance

Analysis

Implementation and Usage

Common Use Cases

Caveats

Summary

- The expression is treated as a list of constant-length expressions, from each possible starting expression to the logical end.
- All of these expressions will differ only around places where alternations were present.
- Therefore, everything leading up to the alternation, and everything after the alternation, will be identical for various expressions.
- We can therefore optimize the expression by refactoring all common prefixes and suffixes.

SCREAM

Rob King

Introduction

The Problem

The Base64

Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
Optimization

Performance
Analysis

Implementation and Usage

Common Use Cases

Caveats

Summary

- In general, the algorithm will run in $O(n^2)$ time, since every addition of an alternation results in twice as many START nodes.
- This can of course become rapidly intractable, since even having 32 alternations would result in over four billion START nodes.
- While we do perform some optimizations to improve the theoretical running time, it is useful to point out that in the vast majority of practical runs, expressions generally have fewer than twenty alternations.

SCREAM

Rob King

Introduction

The Problem

The Base64

Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression

Optimization

Performance

Analysis

Implementation and Usage

Common Use Cases

Caveats

Summary

- In general, the algorithm will run in $O(n^2)$ time, since every addition of an alternation results in twice as many START nodes.
- This can of course become rapidly intractable, since even having 32 alternations would result in over four billion START nodes.
- While we do perform some optimizations to improve the theoretical running time, it is useful to point out that in the vast majority of practical runs, expressions generally have fewer than twenty alternations.

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
OptimizationPerformance
AnalysisImplementation
and Usage

Common Use Cases

Caveats

Summary

- In general, the algorithm will run in $O(n^2)$ time, since every addition of an alternation results in twice as many START nodes.
- This can of course become rapidly intractable, since even having 32 alternations would result in over four billion START nodes.
- While we do perform some optimizations to improve the theoretical running time, it is useful to point out that in the vast majority of practical runs, expressions generally have fewer than twenty alternations.

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
OptimizationPerformance
AnalysisImplementation
and Usage

Common Use Cases

Caveats

Summary

- In general, the algorithm will run in $O(n^2)$ time, since every addition of an alternation results in twice as many START nodes.
- This can of course become rapidly intractable, since even having 32 alternations would result in over four billion START nodes.
- While we do perform some optimizations to improve the theoretical running time, it is useful to point out that in the vast majority of practical runs, expressions generally have fewer than twenty alternations.

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
OptimizationPerformance
AnalysisImplementation
and Usage

Common Use Cases

Caveats

Summary

- The primary space optimization is the shared tail of the core data structure. Outside of pathological cases, this saves considerable storage.
- The alternation processing portion of the algorithm in reality only splits the expression when the two branches of the expression are of unequal length. This reduces the number of paths in many situations.
- Most of the calculations are memoized; that is, they are run only once for any given input for any given position and offset.
- With all of these optimizations taken into account, for most regular expressions in our test data set, runtimes of under three minutes are found (though memory usage is roughly quadratic on the number of operations in the expression).

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
OptimizationPerformance
AnalysisImplementation
and Usage

Common Use Cases

Caveats

Summary

- The primary space optimization is the shared tail of the core data structure. Outside of pathological cases, this saves considerable storage.
- The alternation processing portion of the algorithm in reality only splits the expression when the two branches of the expression are of unequal length. This reduces the number of paths in many situations.
- Most of the calculations are memoized; that is, they are run only once for any given input for any given position and offset.
- With all of these optimizations taken into account, for most regular expressions in our test data set, runtimes of under three minutes are found (though memory usage is roughly quadratic on the number of operations in the expression).

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
OptimizationPerformance
AnalysisImplementation
and Usage

Common Use Cases

Caveats

Summary

- The primary space optimization is the shared tail of the core data structure. Outside of pathological cases, this saves considerable storage.
- The alternation processing portion of the algorithm in reality only splits the expression when the two branches of the expression are of unequal length. This reduces the number of paths in many situations.
- Most of the calculations are memoized; that is, they are run only once for any given input for any given position and offset.
- With all of these optimizations taken into account, for most regular expressions in our test data set, runtimes of under three minutes are found (though memory usage is roughly quadratic on the number of operations in the expression).

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
OptimizationPerformance
AnalysisImplementation
and Usage

Common Use Cases

Caveats

Summary

- The primary space optimization is the shared tail of the core data structure. Outside of pathological cases, this saves considerable storage.
- The alternation processing portion of the algorithm in reality only splits the expression when the two branches of the expression are of unequal length. This reduces the number of paths in many situations.
- Most of the calculations are memoized; that is, they are run only once for any given input for any given position and offset.
- With all of these optimizations taken into account, for most regular expressions in our test data set, runtimes of under three minutes are found (though memory usage is roughly quadratic on the number of operations in the expression).

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
OptimizationPerformance
AnalysisImplementation
and Usage

Common Use Cases

Caveats

Summary

- The primary space optimization is the shared tail of the core data structure. Outside of pathological cases, this saves considerable storage.
- The alternation processing portion of the algorithm in reality only splits the expression when the two branches of the expression are of unequal length. This reduces the number of paths in many situations.
- Most of the calculations are memoized; that is, they are run only once for any given input for any given position and offset.
- With all of these optimizations taken into account, for most regular expressions in our test data set, runtimes of under three minutes are found (though memory usage is roughly quadratic on the number of operations in the expression).

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
OptimizationPerformance
AnalysisImplementation
and Usage

Common Use Cases

Caveats

Summary

- The initial version of the tool was written in Lua.
- However, the tool was quickly rewritten in Erlang.
- Erlang provided a much cleaner method of implementing shared-tail lists, memoization, and bit manipulation.
 - Total implementation is around 1500 lines of heavily-commented Erlang.

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
OptimizationPerformance
AnalysisImplementation
and Usage

Common Use Cases

Caveats

Summary

- The initial version of the tool was written in Lua.
- However, the tool was quickly rewritten in Erlang.
- Erlang provided a much cleaner method of implementing shared-tail lists, memoization, and bit manipulation.
 - Total implementation is around 1500 lines of heavily-commented Erlang.

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
OptimizationPerformance
AnalysisImplementation
and Usage

Common Use Cases

Caveats

Summary

- The initial version of the tool was written in Lua.
- However, the tool was quickly rewritten in Erlang.
- Erlang provided a much cleaner method of implementing shared-tail lists, memoization, and bit manipulation.
 - Total implementation is around 1500 lines of heavily-commented Erlang.

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
OptimizationPerformance
AnalysisImplementation
and Usage

Common Use Cases

Caveats

Summary

- The initial version of the tool was written in Lua.
- However, the tool was quickly rewritten in Erlang.
- Erlang provided a much cleaner method of implementing shared-tail lists, memoization, and bit manipulation.
 - Total implementation is around 1500 lines of heavily-commented Erlang.

SCREAM

Rob King

Introduction

The Problem

The Base64

Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression

Optimization

Performance

Analysis

Implementation
and Usage

Common Use Cases

Caveats

Summary

- The tool's functionality can be divided into the following stages, each of which is well-served by some part of Erlang or the OTP:
 - Parse the regular expression (YECC).
 - Create a series of bitstrings and arrays of bitstrings (Erlang's native bitstring functionality).
 - Create a graph of bitstring nodes, and manipulate the graph (Erlang's DIGRAPH module).
 - Determine unique paths through the graph (DIGRAPH).
 - Compute the transformations on each path, in parallel (Erlang's extremely easy process creation).
 - Combine the results and optimize them (optimization problem was very easy to express in a functional way).
 - Serialize everything out (uhm...IO:FORMAT?).

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
Optimization

Performance
Analysis

Implementation and Usage

Common Use Cases

Caveats

Summary

- Erlang makes it very easy to parallelize computation across multiple threads, processors, and machines.
- The algorithm itself is very amenable to parallelization, and an experimental version of the tool has shown significant speedups.

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
Optimization

Performance
Analysis

Implementation and Usage

Common Use Cases

Caveats

Summary

- Erlang makes it very easy to parallelize computation across multiple threads, processors, and machines.
- The algorithm itself is very amenable to parallelization, and an experimental version of the tool has shown significant speedups.

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
Optimization

Performance
Analysis

Implementation and Usage

Common Use Cases

Caveats

Summary

- While a traditional store-and-forward email server is of course optimal for filtering email, b64re-based signatures on an IPS have been useful in prefiltering email.
- One common use case is filtering obviously malicious email, based on shellcode detection or other parameters, lowering the load on the backend mail server.
- Another interesting application is protecting store-and-forward scanners from themselves. Some systems have bugs that can be triggered by specially formatted emails; an inline IPS can filter these attacks.

SCREAM

Rob King

Introduction

The Problem

The Base64

Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression

Optimization

Performance

Analysis

Implementation
and Usage

Common Use Cases

Caveats

Summary

- While a traditional store-and-forward email server is of course optimal for filtering email, b64re-based signatures on an IPS have been useful in prefiltering email.
- One common use case is filtering obviously malicious email, based on shellcode detection or other parameters, lowering the load on the backend mail server.
- Another interesting application is protecting store-and-forward scanners from themselves. Some systems have bugs that can be triggered by specially formatted emails; an inline IPS can filter these attacks.

SCREAM

Rob King

Introduction

The Problem

The Base64

Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression

Optimization

Performance

Analysis

Implementation
and Usage

Common Use Cases

Caveats

Summary

- While a traditional store-and-forward email server is of course optimal for filtering email, b64re-based signatures on an IPS have been useful in prefiltering email.
- One common use case is filtering obviously malicious email, based on shellcode detection or other parameters, lowering the load on the backend mail server.
- Another interesting application is protecting store-and-forward scanners from themselves. Some systems have bugs that can be triggered by specially formatted emails; an inline IPS can filter these attacks.

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
OptimizationPerformance
AnalysisImplementation
and Usage

Common Use Cases

Caveats

Summary

- While a traditional store-and-forward email server is of course optimal for filtering email, b64re-based signatures on an IPS have been useful in prefiltering email.
- One common use case is filtering obviously malicious email, based on shellcode detection or other parameters, lowering the load on the backend mail server.
- Another interesting application is protecting store-and-forward scanners from themselves. Some systems have bugs that can be triggered by specially formatted emails; an inline IPS can filter these attacks.

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
Optimization

Performance
Analysis

Implementation and Usage

Common Use Cases

Caveats

Summary

- Email is a common source of sensitive data leakage.
- An inline IPS running signatures using this tool can be used to scan for outbound sensitive information, even if an external mail server is used.

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
Optimization

Performance
Analysis

Implementation and Usage

Common Use Cases

Caveats

Summary

- **Email is a common source of sensitive data leakage.**
- An inline IPS running signatures using this tool can be used to scan for outbound sensitive information, even if an external mail server is used.

SCREAM

Rob King

Introduction

The Problem

The Base64

Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression

Optimization

Performance

Analysis

Implementation and Usage

Common Use Cases

Caveats

Summary

- Email is a common source of sensitive data leakage.
- An inline IPS running signatures using this tool can be used to scan for outbound sensitive information, even if an external mail server is used.

SCREAM

Rob King

Introduction

The Problem

The Base64

Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression

Optimization

Performance

Analysis

Implementation and Usage

Common Use Cases

Caveats

Summary

- Expressions processed in this way are inherently looser than their unencoded counterparts. Why?
- Recall that the expression is encoded three times, to account for varying offsets from the beginning of input.
- The expression designed to match when its offset is $n + 1$ from the beginning of the input could match input that is actually at a different modular offset.
- This is mitigated somewhat if the expression can be anchored by long constant strings, or if the beginning of input is known.

SCREAM

Rob King

Introduction

The Problem

The Base64

Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression

Optimization

Performance

Analysis

Implementation and Usage

Common Use Cases

Caveats

Summary

- Expressions processed in this way are inherently looser than their unencoded counterparts. Why?
- Recall that the expression is encoded three times, to account for varying offsets from the beginning of input.
- The expression designed to match when its offset is $n + 1$ from the beginning of the input could match input that is actually at a different modular offset.
- This is mitigated somewhat if the expression can be anchored by long constant strings, or if the beginning of input is known.

SCREAM

Rob King

Introduction

The Problem

The Base64

Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression

Optimization

Performance

Analysis

Implementation and Usage

Common Use Cases

Caveats

Summary

- Expressions processed in this way are inherently looser than their unencoded counterparts. Why?
- Recall that the expression is encoded three times, to account for varying offsets from the beginning of input.
- The expression designed to match when its offset is $n + 1$ from the beginning of the input could match input that is actually at a different modular offset.
- This is mitigated somewhat if the expression can be anchored by long constant strings, or if the beginning of input is known.

SCREAM

Rob King

Introduction

The Problem

The Base64

Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression

Optimization

Performance

Analysis

Implementation and Usage

Common Use Cases

Caveats

Summary

- Expressions processed in this way are inherently looser than their unencoded counterparts. Why?
- Recall that the expression is encoded three times, to account for varying offsets from the beginning of input.
- The expression designed to match when its offset is $n + 1$ from the beginning of the input could match input that is actually at a different modular offset.
- This is mitigated somewhat if the expression can be anchored by long constant strings, or if the beginning of input is known.

SCREAM

Rob King

Introduction

The Problem

The Base64

Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression

Optimization

Performance

Analysis

Implementation and Usage

Common Use Cases

Caveats

Summary

- Expressions processed in this way are inherently looser than their unencoded counterparts. Why?
- Recall that the expression is encoded three times, to account for varying offsets from the beginning of input.
- The expression designed to match when its offset is $n + 1$ from the beginning of the input could match input that is actually at a different modular offset.
- This is mitigated somewhat if the expression can be anchored by long constant strings, or if the beginning of input is known.

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
Optimization

Performance
Analysis

Implementation and Usage

Common Use Cases

Caveats

Summary

- Most Base64 processors allow the insertion of newlines at arbitrary points of the input.
- The encoded expression can be modified to include an optional newline after every character, though this is expensive.
- Some regular expression engines offer the ability to strip newlines from input, removing the burden on the expression.

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
Optimization

Performance
Analysis

Implementation and Usage

Common Use Cases

Caveats

Summary

- Most Base64 processors allow the insertion of newlines at arbitrary points of the input.
- The encoded expression can be modified to include an optional newline after every character, though this is expensive.
- Some regular expression engines offer the ability to strip newlines from input, removing the burden on the expression.

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
Optimization
Performance
AnalysisImplementation
and Usage

Common Use Cases

Caveats

Summary

- Most Base64 processors allow the insertion of newlines at arbitrary points of the input.
- The encoded expression can be modified to include an optional newline after every character, though this is expensive.
- Some regular expression engines offer the ability to strip newlines from input, removing the burden on the expression.

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
OptimizationPerformance
AnalysisImplementation
and Usage

Common Use Cases

Caveats

Summary

- Most Base64 processors allow the insertion of newlines at arbitrary points of the input.
- The encoded expression can be modified to include an optional newline after every character, though this is expensive.
- Some regular expression engines offer the ability to strip newlines from input, removing the burden on the expression.

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
OptimizationPerformance
AnalysisImplementation
and UsageCommon Use Cases
Caveats

Summary

- Using this tool, users can reap several benefits:
 - Pre-filter encoded data streams, lessening the need for expensive decoding.
 - Protect infrastructure that provides decoding services but is itself vulnerable
 - Inspect encoded streams within incurring the overhead of decoding.
- This talk illustrates some interesting challenges when inspecting encoded streams of data, especially in a performance-critical way, and gives some possible solutions.
- This talk illustrates an interesting application of Erlang outside of its normal domain of highly parallel soft real-time applications, and shows how its strengths make it well-suited for this sort of problem.
 - This tool served as a “back door” for getting Erlang into DVLabs.

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
OptimizationPerformance
AnalysisImplementation
and UsageCommon Use Cases
Caveats

Summary

- Using this tool, users can reap several benefits:
 - Pre-filter encoded data streams, lessening the need for expensive decoding.
 - Protect infrastructure that provides decoding services but is itself vulnerable
 - Inspect encoded streams within incurring the overhead of decoding.
- This talk illustrates some interesting challenges when inspecting encoded streams of data, especially in a performance-critical way, and gives some possible solutions.
- This talk illustrates an interesting application of Erlang outside of its normal domain of highly parallel soft real-time applications, and shows how its strengths make it well-suited for this sort of problem.
 - This tool served as a “back door” for getting Erlang into DVLabs.

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
OptimizationPerformance
AnalysisImplementation
and UsageCommon Use Cases
Caveats

Summary

- Using this tool, users can reap several benefits:
 - Pre-filter encoded data streams, lessening the need for expensive decoding.
 - Protect infrastructure that provides decoding services but is itself vulnerable
 - Inspect encoded streams within incurring the overhead of decoding.
- This talk illustrates some interesting challenges when inspecting encoded streams of data, especially in a performance-critical way, and gives some possible solutions.
- This talk illustrates an interesting application of Erlang outside of its normal domain of highly parallel soft real-time applications, and shows how its strengths make it well-suited for this sort of problem.
 - This tool served as a “back door” for getting Erlang into DVLabs.

SCREAM

Rob King

Introduction

The Problem

The Base64

Algorithm

Regular Expressions

The Algorithm

Ways of Solving the

Problem

Encoding Operations

Performance

Expression

Optimization

Performance

Analysis

Implementation

and Usage

Common Use Cases

Caveats

Summary

- Using this tool, users can reap several benefits:
 - Pre-filter encoded data streams, lessening the need for expensive decoding.
 - Protect infrastructure that provides decoding services but is itself vulnerable
 - Inspect encoded streams within incurring the overhead of decoding.
- This talk illustrates some interesting challenges when inspecting encoded streams of data, especially in a performance-critical way, and gives some possible solutions.
- This talk illustrates an interesting application of Erlang outside of its normal domain of highly parallel soft real-time applications, and shows how its strengths make it well-suited for this sort of problem.
 - This tool served as a “back door” for getting Erlang into DVLabs.

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
OptimizationPerformance
AnalysisImplementation
and UsageCommon Use Cases
Caveats

Summary

- Using this tool, users can reap several benefits:
 - Pre-filter encoded data streams, lessening the need for expensive decoding.
 - Protect infrastructure that provides decoding services but is itself vulnerable
 - Inspect encoded streams within incurring the overhead of decoding.
- This talk illustrates some interesting challenges when inspecting encoded streams of data, especially in a performance-critical way, and gives some possible solutions.
- This talk illustrates an interesting application of Erlang outside of its normal domain of highly parallel soft real-time applications, and shows how its strengths make it well-suited for this sort of problem.
 - This tool served as a “back door” for getting Erlang into DVLabs.

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
OptimizationPerformance
AnalysisImplementation
and UsageCommon Use Cases
Caveats

Summary

- Initially, Erlang suffered from a “not invented here” type mentality, and a fear that “no one knows Erlang”.
- This tool really couldn't have been written to run efficiently very easily in any other language.
- Thanks to this tool, we've encouraged adoption of Erlang in various other portions of the enterprise:
 - A compiler written in Erlang serves as the first stage of our build process.
 - Erlang is now being integrated into various testing tools that we use for quality assurance and threat analysis.
 - Interesting tools used for static analysis of binaries are being written in Erlang (maybe something to talk about at Erlang Factory 2011?)

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
OptimizationPerformance
AnalysisImplementation
and UsageCommon Use Cases
Caveats

Summary

- Initially, Erlang suffered from a “not invented here” type mentality, and a fear that “no one knows Erlang”.
- This tool really couldn't have been written to run efficiently very easily in any other language.
- Thanks to this tool, we've encouraged adoption of Erlang in various other portions of the enterprise:
 - A compiler written in Erlang serves as the first stage of our build process.
 - Erlang is now being integrated into various testing tools that we use for quality assurance and threat analysis.
 - Interesting tools used for static analysis of binaries are being written in Erlang (maybe something to talk about at Erlang Factory 2011?)

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
OptimizationPerformance
AnalysisImplementation
and Usage

Common Use Cases

Caveats

Summary

- Initially, Erlang suffered from a “not invented here” type mentality, and a fear that “no one knows Erlang”.
- This tool really couldn't have been written to run efficiently very easily in any other language.
- Thanks to this tool, we've encouraged adoption of Erlang in various other portions of the enterprise:
 - A compiler written in Erlang serves as the first stage of our build process.
 - Erlang is now being integrated into various testing tools that we use for quality assurance and threat analysis.
 - Interesting tools used for static analysis of binaries are being written in Erlang (maybe something to talk about at Erlang Factory 2011?)

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
OptimizationPerformance
AnalysisImplementation
and Usage

Common Use Cases

Caveats

Summary

- Initially, Erlang suffered from a “not invented here” type mentality, and a fear that “no one knows Erlang”.
- This tool really couldn't have been written to run efficiently very easily in any other language.
- Thanks to this tool, we've encouraged adoption of Erlang in various other portions of the enterprise:
 - A compiler written in Erlang serves as the first stage of our build process.
 - Erlang is now being integrated into various testing tools that we use for quality assurance and threat analysis.
 - Interesting tools used for static analysis of binaries are being written in Erlang (maybe something to talk about at Erlang Factory 2011?)

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
OptimizationPerformance
AnalysisImplementation
and Usage

Common Use Cases

Caveats

Summary

- Initially, Erlang suffered from a “not invented here” type mentality, and a fear that “no one knows Erlang”.
- This tool really couldn’t have been written to run efficiently very easily in any other language.
- Thanks to this tool, we’ve encouraged adoption of Erlang in various other portions of the enterprise:
 - A compiler written in Erlang serves as the first stage of our build process.
 - Erlang is now being integrated into various testing tools that we use for quality assurance and threat analysis.
 - Interesting tools used for static analysis of binaries are being written in Erlang (maybe something to talk about at Erlang Factory 2011?)

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
OptimizationPerformance
AnalysisImplementation
and Usage

Common Use Cases

Caveats

Summary

- Initially, Erlang suffered from a “not invented here” type mentality, and a fear that “no one knows Erlang”.
- This tool really couldn’t have been written to run efficiently very easily in any other language.
- Thanks to this tool, we’ve encouraged adoption of Erlang in various other portions of the enterprise:
 - A compiler written in Erlang serves as the first stage of our build process.
 - Erlang is now being integrated into various testing tools that we use for quality assurance and threat analysis.
 - Interesting tools used for static analysis of binaries are being written in Erlang (maybe something to talk about at Erlang Factory 2011?)

SCREAM

Rob King

Introduction

The Problem

The Base64
Algorithm

Regular Expressions

The Algorithm

Ways of Solving the
Problem

Encoding Operations

Performance

Expression
OptimizationPerformance
AnalysisImplementation
and Usage

Common Use Cases

Caveats

Summary

- Initially, Erlang suffered from a “not invented here” type mentality, and a fear that “no one knows Erlang”.
- This tool really couldn’t have been written to run efficiently very easily in any other language.
- Thanks to this tool, we’ve encouraged adoption of Erlang in various other portions of the enterprise:
 - A compiler written in Erlang serves as the first stage of our build process.
 - Erlang is now being integrated into various testing tools that we use for quality assurance and threat analysis.
 - Interesting tools used for static analysis of binaries are being written in Erlang (maybe something to talk about at Erlang Factory 2011?)

SCREAM

Rob King

Introduction

The Problem
The Base64
Algorithm
Regular Expressions

The Algorithm

Ways of Solving the
Problem
Encoding Operations

Performance

Expression
Optimization
Performance
Analysis

Implementation
and Usage

Common Use Cases
Caveats

Summary

- Rob at Work: rking@tippingpoint.com
- Rob at Home: jking@deadpixi.com