An Applicative Application

Application

How to obtain OpenTSDB query batching by exploiting the applicative structure of a computation.

A Word of Warning

A Word of Warning I had trouble explaining this stuff



RE,

Rob Norris @tpolecat

Sorry, I need to focus on something else right now. I'm sure someone can help you work through it.



Oct 04 2017 02:24

Ionuț G. Stan @igstan

No worries, I like banging my head against this wall actually :)



Ionuț G. Stan @igstan	Oct 04 2017 03:06
Well, it doesn't have to return, it's an applicative, so its continuation can be	
saved somewhere, right?	
Fabio Labella @SystemFw	Oct 04 2017 03:06



Fabio Labella @SystemFw	
I can't parse that	

Fabio Labella @SystemFw I'm busy today	Oct 04 2017 13:46
Ionuț G. Stan @igstan No worries, you've helped a lot anyways.	Oct 04 2017 13:46
Fabio Labella @SystemFw I'll try to show actual code tonight or tomorrow (UK time)	Oct 04 2017 13:46

• The Problem

- The Problem
 - OpenTSDB Overview Time Series Data Storage

- The Problem
 - OpenTSDB Overview Time Series Data Storage
 - Business Use Case Time Series Data Aggregation

- The Problem
 - OpenTSDB Overview Time Series Data Storage
 - Business Use Case Time Series Data Aggregation
- The Solution

- The Problem
 - OpenTSDB Overview Time Series Data Storage
 - Business Use Case Time Series Data Aggregation
- The Solution
 - Inspiration

- The Problem
 - OpenTSDB Overview Time Series Data Storage
 - Business Use Case Time Series Data Aggregation
- The Solution
 - Inspiration
 - Intuitions Functions, Monads & Applicatives

- The Problem
 - OpenTSDB Overview Time Series Data Storage
 - Business Use Case Time Series Data Aggregation
- The Solution
 - Inspiration
 - Intuitions Functions, Monads & Applicatives
 - Code

- The Problem
 - OpenTSDB Overview Time Series Data Storage
 - Business Use Case Time Series Data Aggregation
- The Solution
 - Inspiration
 - Intuitions Functions, Monads & Applicatives
 - Code
- Future Work

OpenTSDB Overview Time Series Data Storage

```
http POST 'localhost:4242/api/put' << EOF</pre>
L
  {
    "metric": "metric-a",
    "timestamp": 1538092800,
    "value": 10,
    "tags": {
     "t1": "v1",
      "t2": "v2"
    }
  },
{
    "metric": "metric-b",
    "timestamp": 1538092801,
    "value": 20,
    "tags": {
     "t1": "v1",
      "t2": "v2"
    }
  }
EOF
```

```
http POST 'localhost:4242/api/put' << EOF</pre>
L
  {
    "metric": "metric-a",
    "timestamp": 1538092800,
    "value": 10,
    "tags": {
      "t1": "v1",
      "t2": "v2"
    }
  },
{
    "metric": "metric-b",
    "timestamp": 1538092801,
    "value": 20,
    "tags": {
     "t1": "v1",
      "t2": "v2"
    }
  }
EOF
```

```
http POST 'localhost:4242/api/put' << EOF</pre>
L
  {
    "metric": "metric-a",
    "timestamp": 1538092800,
    "value": 10,
    "tags": {
      "t1": "v1",
      "t2": "v2"
    }
  },
{
    "metric": "metric-b",
    "timestamp": 1538092801,
    "value": 20,
    "tags": {
     "t1": "v1",
      "t2": "v2"
    }
  }
EOF
```

```
http POST 'localhost:4242/api/put' << EOF</pre>
L
  {
    "metric": "metric-a",
    "timestamp": 1538092800,
    "value": 10,
    "tags": {
      "t1": "v1",
      "t2": "v2"
    }
  },
{
    "metric": "metric-b",
    "timestamp": 1538092801,
    "value": 20,
    "tags": {
     "t1": "v1",
      "t2": "v2"
    }
  }
EOF
```

```
http POST 'localhost:4242/api/put' << EOF</pre>
L
  {
    "metric": "metric-a",
    "timestamp": 1538092800,
    "value": 10,
    "tags": {
      "t1": "v1",
      "t2": "v2"
    }
  },
{
    "metric": "metric-b",
    "timestamp": 1538092801,
    "value": 20,
    "tags": {
     "t1": "v1",
      "t2": "v2"
    }
  }
EOF
```

```
http POST 'localhost:4242/api/query' << EOF</pre>
{
  "start": "2018/05/01 00:10:00",
  "end": "2018/05/01 00:10:05",
  "timezone": "UTC",
  "queries": [
    {
      "metric": "metric-a",
      "aggregator": "none",
      "tags": {}
    },
{
      "metric": "metric-b",
      "aggregator": "none",
      "tags": {}
    }
  ]
}
EOF
```

			localhost	×	1 0 +
	OPENTSDB				
Grap	oh Stats Logs Version				
From	To (now)		WxH: 1420x645 Global annotations		
me	etric-a metric-b +	JU.21.00	Axes Key Style		
Metr	ric: metric-a Ra	Rate Rate Ctr Right Axis	Y Y2 Label		
Tage	s Ra	te Ctr Reset: 0	Format		
		gregator: sum 📀 Downsample	Range [90:12] Log scale		
Cache b	a hit (disk) 1200 points retrieved 1200 points plotted in 17ms	vg v 10m none 🖓]		
120		· 1			metric-a(tag-a=val-a, tag-b=val-b)
					me <u>tric-b(tag-a-vai-a, tag-b-vai-b) ——</u>
115 -					
110					
	+++ **	ĸ x ¥ tt +X ¥ + X ¥X+ ¥XX ₩ ₽N ₩ ₩ K X H+1 ₩₩₽14 X ₩	++ ++×++3##¥ × + × +++ ++ ++ ++ ++ ++ ₩		
105		1999年1991年1999年1997年1999年1997年1997年1997	· \$4 \$4 \$2###**********************************	- ALE	
103	2011 - 2010 - 2017 - 1000 - 2010 - 2017 - 2017 2011 - 2017 -	- ^ ^ ^ ^ ^ ^ ^ ^ ^ ^ ^ ^ ^ ^ ^ ^ ^ ^ ^		an a	S (17) (17) (17) Kan Anno (1) (17) (18) (27) (27) (20) (17) + MANE (18) (27) (27) (27) (17)
	- k +++ ++++++ +++++ +++++ ++++++	×++-) = ki ist+set+ +++ ki+seins: ×++-×>>×× × × × × +× +=	34 34 14 14 (34 44 44 14 14 14 14 14 14 14 14 14 14 14	\$\$````	
100 -		┄┼┄┼ ╎┼ ╱┼ ┙ ╫┈┼╈╈┼╢╳╶┼┤╺╫┤╬╣╴	╘╶┽┄╶╳╌╳╌╳┾ ╫ ┄┄╌╳┝╱┥╌╏┼╫┾┝╌╌╌╌╫┝╬╸╌╌╳┝┿┥	╋┼┈┈┾╫┿┿	₽····×₩·↓··∔x·×₩↓·····↓

Business Use Case Time Series Data Aggregation

• We store historical data in OpenTSDB

- We store historical data in OpenTSDB
- Basis for end-of-month reports

- We store historical data in OpenTSDB
- Basis for end-of-month reports
- Past: everything computed at end of month from raw OpenTSDB data

- We store historical data in OpenTSDB
- Basis for end-of-month reports
- Past: everything computed at end of month from raw OpenTSDB data
- Now: pre-compute aggregations continuously at various frequencies, e.g., 5 minutes, and use these for reporting

- We store historical data in OpenTSDB
- Basis for end-of-month reports
- Past: everything computed at end of month from raw OpenTSDB data
- Now: pre-compute aggregations continuously at various frequencies, e.g., 5 minutes, and use these for reporting
- We call these aggregations roll-ups

- We store historical data in OpenTSDB
- Basis for end-of-month reports
- Past: everything computed at end of month from raw OpenTSDB data
- Now: pre-compute aggregations continuously at various frequencies, e.g., 5 minutes, and use these for reporting
- We call these aggregations roll-ups
- Decreases time to generate reports

Roll-Ups as Functions

def rollup(metric: List[Datapoint]): Double

```
case class Datapoint(
   timestamp: Instant,
   value: Double,
)
```

def rollup(metric: List[Datapoint]): Double

```
case class Datapoint(
   timestamp: Instant,
   value: Double,
)
```

```
def rollup(
   metric0: List[Datapoint],
   metric1: List[Datapoint],
   ...
   metricN: List[Datapoint],
): Double
```

Goal: Optimize Network Access

Inspiration
GitHu	ubGist Search All gists GitHub	New gist
chr Last	nris-taylor / IOAction.hs st active 2 years ago • Report abuse	★ Star 6 % Fork 1
<> Coo	ode Revisions 2 🖈 Stars 6 🖗 Forks 1	Embed - <script \="" `seqio`="" is="" name?"="" src="https://gist.</th></tr><tr><th>Code for</th><th>or my blog post about pure I/O</th><th></th></tr><tr><td>•• IOAc</td><td>Action.hs</td><td>Raw</td></tr><tr><th>1</th><th>data IOAction a = Return a</th><th></th></tr><tr><td>2</td><td>Put String (IOAction a)</td><td></td></tr><tr><td>3</td><td>Get (String -> IOAction a)</td><td></td></tr><tr><td>4</td><td></td><td></td></tr><tr><td>5</td><td>get = Get Return</td><td></td></tr><tr><td>6</td><td><pre>put s = Put s (Return ())</pre></td><td></td></tr><tr><td>7</td><td></td><td></td></tr><tr><td>8</td><td><pre>seqio :: IOAction a -> (a -> IOAction b) -> IOAction b</pre></td><td></td></tr><tr><td>9</td><td>seqio (Return a) f = f a</td><td></td></tr><tr><td>10</td><td>seqio (Put s io) f = Put s (seqio io f)</td><td></td></tr><tr><td>11</td><td><pre>seqio (Get g) f = Get (\s -> seqio (g s) f)</pre></td><td></td></tr><tr><td>12</td><td></td><td></td></tr><tr><td>13</td><td>echo = get `seqio` put</td><td></td></tr><tr><td>14</td><td></td><td></td></tr><tr><td>15</td><td>hello = put " what="" your=""></script>

```
✓ IOAction.hs
```

```
data IOAction a = Return a
 1
 2
                     Put String (IOAction a)
                     Get (String -> IOAction a)
 3
 4
 5
      get = Get Return
      put s = Put s (Return ())
 6
 7
      seqio :: IOAction a -> (a -> IOAction b) -> IOAction b
 8
      seqio (Return a) f = f a
 9
      seqio (Put s io) f = Put s (seqio io f)
10
      seqio (Get g) f = Get (\s -> seqio (g s) f)
11
12
      echo = get `seqio` put
13
14
      hello = put "What is your name?" `seqio` \_ ->
15
                                          `seqio` \name ->
16
             get
17
             put "What is your age?"
                                          `segio` \ ->
                                          `seqio` \age ->
18
             get
             put ("Hello " ++ name ++ "!") `segio` \ ->
19
             put ("You are " ++ age ++ " years old")
20
21
22
      hello2 = do put "What is your name?"
23
                 name <- get
                 put "What is your age?"
24
25
                 age <- get
```

stack overflow

166

Search...

What are free monads?





don't understand any of this Haskell mumbo – I--"""------"" Aug 13 at 11:03



Move Over Free Monads: Make Way for Free Applicatives

— John de Goes





CC

Free Applicative Functors

Paolo Capriotti

Functional Programming Laboratory University of Nottingham pvc@cs.nott.ac.uk Ambrus Kaposi

Functional Programming Laboratory University of Nottingham auk@cs.nott.ac.uk

Abstract

Applicative functors ([9]) are a generalisation of monads. Both allow expressing effectful computations into an otherwise pure language, like Haskell ([8]).

Applicative functors are to be preferred to monads when the structure of a computation is fixed *a priori*. That makes it possible to perform certain kinds of static analysis on applicative values.

We define a notion of *free applicative functor*, prove that it satisfies the appropriate laws, and that the construction is left adjoint to a suitable forgetful functor.

We show how free applicative functors can be used to implement embedded DSLs which can be statically analysed.

Categories and Subject Descriptors D.3.3 [*Programming Languages*]: Language Constructs and Features

General Terms Haskell, Free Applicative Functors

Keywords Applicative Functors, Parametricity, Adjoints

1. Introduction

Free monads in Haskell are a very well-known and practically used construction. Given any endofunctor f, the free monad on f is given by a simple inductive definition:

```
data Free f a
    = Return a
    | Free (f (Free f a))
```

The typical use case for this construction is creating embedded

- We give two definitions of *free applicative functor* in Haskell (section 2), and show that they are equivalent (section 5).
- We prove that our definition is correct, in the sense that it really is an applicative functor (section 6), and that it is "free" in a precise sense (section 7).
- We present a number of examples where the use of free applicative functors helps make the code more elegant, removes duplication or enables certain kinds of optimizations which are not possible when using free monads. We describe the differences between expressivity of DSLs using free applicatives and free monads (section 3).
- We compare our definition to other existing implementations of the same idea (section 9).

This paper is aimed at programmers with a working knowledge of Haskell. Familiarity with applicative functors is not required, although it is helpful to understand the motivation behind this work. We make use of category theoretical concepts to justify our definition, but the Haskell code we present can also stand on its own.

1.1 Applicative functors

Applicative functors (also called *idioms*) were first introduced in [9] as a way to provide a lighter notation for monads. They have since been used in a variety of different applications, including efficient parsing (see section 1.4), regular expressions and bidirectional routing.

Free Applicative Functors

Paolo Capriotti

Functional Programming Laboratory University of Nottingham pvc@cs.nott.ac.uk Ambrus Kaposi

Functional Programming Laboratory University of Nottingham auk@cs.nott.ac.uk

Abstract

Applicative functors ([9]) are a generalisation of monads. Both allow expressing effectful computations into an otherwise pure language, like Haskell ([8]).

Applicative functors are to be preferred to monads when the structure of a computation is fixed *a priori*. That makes it possible to perform certain kinds of static analysis on applicative values.

We define a notion of *free applicative functor*, prove that it satisfies the appropriate laws, and that the construction is left adjoint to a suitable forgetful functor.

We show how free applicative functors can be used to implement embedded DSLs which can be statically analysed.

Categories and Subject Descriptors D.3.3 [*Programming Languages*]: Language Constructs and Features

General Terms Haskell, Free Applicative Functors

Keywords Applicative Functors, Parametricity, Adjoints

1. Introduction

Free monads in Haskell are a very well-known and practically used construction. Given any endofunctor f, the free monad on f is given by a simple inductive definition:

```
data Free f a
    = Return a
    | Free (f (Free f a))
```

The typical use case for this construction is creating embedded

- We give two definitions of *free applicative functor* in Haskell (section 2), and show that they are equivalent (section 5).
- We prove that our definition is correct, in the sense that it really is an applicative functor (section 6), and that it is "free" in a precise sense (section 7).
- We present a number of examples where the use of free applicative functors helps make the code more elegant, removes duplication or enables certain kinds of optimizations which are not possible when using free monads. We describe the differences between expressivity of DSLs using free applicatives and free monads (section 3).
- We compare our definition to other existing implementations of the same idea (section 9).

This paper is aimed at programmers with a working knowledge of Haskell. Familiarity with applicative functors is not required, although it is helpful to understand the motivation behind this work. We make use of category theoretical concepts to justify our definition, but the Haskell code we present can also stand on its own.

1.1 Applicative functors

Applicative functors (also called *idioms*) were first introduced in [9] as a way to provide a lighter notation for monads. They have since been used in a variety of different applications, including efficient parsing (see section 1.4), regular expressions and bidirectional routing.

Abstract

Applicative functors ([9]) are a generalisation of monads. Both allow expressing effectful computations into an otherwise pure language, like Haskell ([8]).

Applicative functors are to be preferred to monads when the structure of a computation is fixed *a priori*. That makes it possible to perform certain kinds of static analysis on applicative values. We define a notion of *free applicative functor*, prove that it satisfies the appropriate laws, and that the construction is left adjoint to a suitable forgetful functor.

We show how free applicative functors can be used to implement embedded DSLs which can be statically analysed.

Categories and Subject Descriptors D.3.3 [*Programming Languages*]: Language Constructs and Features

General Terms Haskell, Free Applicative Functors

Keywords Applicative Functors, Parametricity, Adjoints

1. Introduction

Free monads in Haskell are a very well-known and practically used construction. Given any endofunctor f, the free monad on f is given by a simple inductive definition:

The typical use case for this construction is creating embedded

- We give two d (section 2), and
- We prove that of is an applicative precise sense (set a set a
- We present a r plicative functo duplication or o not possible wi ences between free monads (so
- We compare ou the same idea (

This paper is aime of Haskell. Famili although it is help work. We make us definition, but the own.

1.1 Applicative

Applicative functor as a way to provide been used in a vari parsing (see section ing.

Abstract

Applicative functors ([9]) are a generalisation of monads. Both allow expressing effectful computations into an otherwise pure language, like Haskell ([8]).

Applicative functors are to be preferred to monads when the structure of a computation is fixed *a priori*. That makes it possible to perform certain kinds of static analysis on applicative values. We define a notion of *free applicative functor*, prove that it satisfies the appropriate laws, and that the construction is left adjoint to a suitable forgetful functor.

We show how free applicative functors can be used to implement embedded DSLs which can be statically analysed.

Categories and Subject Descriptors D.3.3 [*Programming Languages*]: Language Constructs and Features

General Terms Haskell, Free Applicative Functors

Keywords Applicative Functors, Parametricity, Adjoints

1. Introduction

Free monads in Haskell are a very well-known and practically used construction. Given any endofunctor f, the free monad on f is given by a simple inductive definition:

The typical use case for this construction is creating embedded

- We give two do (section 2), and
- We prove that of is an applicative precise sense (set a set a
- We present a r plicative functo duplication or e not possible with ences between free monads (se
- We compare ou the same idea (

This paper is aime of Haskell. Famili although it is help work. We make us definition, but the own.

1.1 Applicative

Applicative functor as a way to provide been used in a vari parsing (see section ing. Similar Ideas: Hax

There is no Fork: an Abstraction for Efficient, Concurrent, and Concise Data Access

Simon Marlow Facebook

Facebook smarlow@fb.com Louis Brandy Facebook Idbrandy@fb.com Jonathan Coens Facebook jon.coens@fb.com Jon Purdy Facebook jonp@fb.com

Abstract

We describe a new programming idiom for concurrency, based on Applicative Functors, where concurrency is implicit in the Applicative <*> operator. The result is that concurrent programs can be written in a natural applicative style, and they retain a high degree of clarity and modularity while executing with maximal concurrency. This idiom is particularly useful for programming against external data sources, where the application code is written without the use of explicit concurrency constructs, while the implementation is able to batch together multiple requests for data from the same source, and fetch data from multiple sources concurrently. Our abstraction uses a cache to ensure that multiple requests for the same data return the same result, which frees the programmer from having to arrange to fetch data only once, which in turn leads to greater modularity.

While it is generally applicable, our technique was designed with a particular application in mind: an internal service at Facebook that identifies particular types of content and takes actions based on it. Our application has a large body of business logic that fetches data from several different external sources. The framework described in this paper enables the business logic to execute efficiently by automatically fetching data concurrently; we present some preliminary results.

Keywords Haskell; concurrency; applicative; monad; data-fetching; distributed

1. Introduction

concise business logic, uncluttered by performance-related details. In particular the programmer should not need to be concerned with accessing external data efficiently. However, one particular problem often arises that creates a tension between conciseness and efficiency in this setting: accessing multiple remote data sources efficiently requires *concurrency*, and that normally requires the programmer to intervene and program the concurrency explicitly.

When the business logic is only concerned with *reading* data from external sources and not *writing*, the programmer doesn't care about the order in which data accesses happen, since there are no side-effects that could make the result different when the order changes. So in this case the programmer would be entirely happy with not having to specify either ordering or concurrency, and letting the system perform data access in the most efficient way possible. In this paper we present an embedded domain-specific language (EDSL), written in Haskell, that facilitates this style of programming, while automatically extracting and exploiting any concurrency inherent in the program.

Our contributions can be summarised as follows:

- We present an Applicative abstraction that allows implicit concurrency to be extracted from computations written with a combination of Monad and Applicative. This is an extension of the idea of concurrency monads [10], using Applicative <*> as a way to introduce concurrency (Section 4). We then develop the idea into an abstraction that supports concurrent access to remote data (Section 5), and failure (Section 8).
- We show how to add a *cache* to the framework (Section 6).

There is no Fork: an Abstraction for Efficient, Concurrent, and Concise Data Access

Simon Marlow Facebook

smarlow@fb.com

Louis Brandy Facebook Idbrandy@fb.com Jonathan Coens Facebook jon.coens@fb.com Jon Purdy Facebook jonp@fb.com

Abstract

We describe a new programming idiom for concurrency, based on Applicative Functors, where concurrency is implicit in the Applicative <*> operator. The result is that concurrent programs can be written in a natural applicative style, and they retain a high degree of clarity and modularity while executing with maximal concurrency. This idiom is particularly useful for programming against external data sources, where the application code is written without the use of explicit concurrency constructs, while the implementation is able to batch together multiple requests for data from the same source, and fetch data from multiple sources concurrently. Our abstraction uses a cache to ensure that multiple requests for the same data return the same result, which frees the programmer from having to arrange to fetch data only once, which in turn leads to greater modularity.

While it is generally applicable, our technique was designed with a particular application in mind: an internal service at Facebook that identifies particular types of content and takes actions based on it. Our application has a large body of business logic that fetches data from several different external sources. The framework described in this paper enables the business logic to execute efficiently by automatically fetching data concurrently; we present some preliminary results.

Keywords Haskell; concurrency; applicative; monad; data-fetching; distributed

1. Introduction

concise business logic, uncluttered by performance-related details. In particular the programmer should not need to be concerned with accessing external data efficiently. However, one particular problem often arises that creates a tension between conciseness and efficiency in this setting: accessing multiple remote data sources efficiently requires *concurrency*, and that normally requires the programmer to intervene and program the concurrency explicitly.

When the business logic is only concerned with *reading* data from external sources and not *writing*, the programmer doesn't care about the order in which data accesses happen, since there are no side-effects that could make the result different when the order changes. So in this case the programmer would be entirely happy with not having to specify either ordering or concurrency, and letting the system perform data access in the most efficient way possible. In this paper we present an embedded domain-specific language (EDSL), written in Haskell, that facilitates this style of programming, while automatically extracting and exploiting any concurrency inherent in the program.

Our contributions can be summarised as follows:

- We present an Applicative abstraction that allows implicit concurrency to be extracted from computations written with a combination of Monad and Applicative. This is an extension of the idea of concurrency monads [10], using Applicative <*> as a way to introduce concurrency (Section 4). We then develop the idea into an abstraction that supports concurrent access to remote data (Section 5), and failure (Section 8).
- We show how to add a *cache* to the framework (Section 6).

Abstract

We describe a new programming idiom for concurrency, based on Applicative Functors, where concurrency is implicit in the Applicative <*> operator. The result is that concurrent programs can be written in a natural applicative style, and they retain a high degree of clarity and modularity while executing with maximal concurrency. This idiom is particularly useful for programming against external data sources, where the application code is written without the use of explicit concurrency constructs, while the implementation is able to batch together multiple requests for data from the same source, and fetch data from multiple sources concurrently. Our abstraction uses a cache to ensure that multiple requests for the same data return the same result, which frees the programmer from having to arrange to fetch data only once, which in turn leads to greater modularity.

While it is generally applicable, our technique was designed with a particular application in mind: an internal service at Facebook that identifies particular types of content and takes actions based on it. Our application has a large body of business logic that fetches data from several different external sources. The framework described in this paper enables the business logic to execute efficiently by automatically fetching data concurrently; we present some preliminary results.

Keywords Haskell; concurrency; applicative; monad; data-fetching; distributed

1. Introduction

concise business log In particular the prewith accessing exter problem often arises efficiency in this set efficiently requires programmer to interv

When the busine from external source care about the order are no side-effects to order changes. So in happy with not having and letting the system possible. In this paper language (EDSL), we programming, while concurrency inherent Our contributions

- We present an A concurrency to b combination of M of the idea of cor as a way to introc the idea into an remote data (Sec
- We show how to

Applicative lets us perform global optimizations on the Abstract Syntax Tree of an Embedded DSL.

Intuitions Functions, Monads & Applicatives

f(a)

def example1[A, B](f: A => B, a: A): B =
 f(a)

```
def example1[A, B](f: A => B, a: A): B =
  f(a)

def example2[F[_]: Applicative, A, B](f: A => B, a: A): F[B] = {
  val ff = Applicative[F].pure(f)
  val fa = Applicative[F].pure(a)
  ff.ap(fa)
}
```

```
def example1[A, B](f: A => B, a: A): B =
  f(a)

def example2[F[_]: Applicative, A, B](f: A => B, a: A): F[B] = {
  val ff = Applicative[F].pure(f)
  val fa = Applicative[F].pure(a)
  ff.ap(fa)
}
def example3[F[_]: Applicative, A, B](f: F[A => B], a: F[A]): F[B] =
  f.ap(a)
```

```
def example1[A, B](f: A => B, a: A): B =
  f.apply(a)

def example2[F[_]: Applicative, A, B](f: A => B, a: A): F[B] = {
  val ff = Applicative[F].pure(f)
  val fa = Applicative[F].pure(a)
  ff.ap(fa)
}
def example3[F[_]: Applicative, A, B](f: F[A => B], a: F[A]): F[B] =
  f.ap(a)
```

Applicative Lets us embed function-like DSLs into our programs.

Data and Effect Dependencies



Functions: 1

Possible Side-Effects Without Data-Dependencies // Potential side-effects, but we don't know.
def foo(): String = ???
def bar(): String = ???

```
// Potential side-effects, but we don't know.
def foo(): String = ???
def bar(): String = ???
```

// No data dependency between function calls.
val a = foo()
val b = bar()

Functions: 2

Possible Side-Effects With Data-Dependencies // Potential side-effects, but we don't know.
def foo(): String = ???
def bar(a: String): String = ???

```
// Potential side-effects, but we don't know.
def foo(): String = ???
def bar(a: String): String = ???
```

// Data-dependency between functions calls.
val a = foo()
val b = bar(a)

Functions: 3

Without Side-Effects Without Data-Dependencies

// Effectful functions; encoded in signature.

- def foo(): Task[String] = ???
- def bar(): Task[String] = ???

// Effectful functions; encoded in signature. def foo(): Task[String] = ??? def bar(): Task[String] = ???

// No data dependency between function calls.
val a = foo()
val b = bar()

```
// Effectful functions; encoded in signature.
def foo(): Task[String] = ???
def bar(): Task[String] = ???
```

```
// No data dependency between function calls.
val a = foo()
val b = bar()
val c = a.flatMap(_ => b) // sequential
```

```
// Effectful functions; encoded in signature.
def foo(): Task[String] = ???
def bar(): Task[String] = ???
```

```
// No data dependency between function calls.
val a = foo()
val b = bar()
val c = a.flatMap(_ => b) // sequential
val d = (a, b).mapN((a, b) => ...)
```

c.unsafeRunSync()

Functions: 4

Without Side-Effects With Data-Dependencies

// Effectful functions; encoded in signature. def foo(): IO[String] = ??? def bar(a: String): IO[String] = ???

// Effectful functions; encoded in signature. def foo(): IO[String] = ??? def bar(a: String): IO[String] = ???

// Data-dependency between functions calls.
val c = foo().flatMap(a => bar(a))
c.unsafeRunSync()
Flashback Applicative lets us embed function-like DSLs into our programs.

Applicative lets us perform global optimizations on the Abstract Syntax Tree of an EDSL.

Roll-Ups are Functions Idea: come up with an EDSL to model time series

data fetching and constrain its usage to an Applicative interface.

Then, statically analyze the EDSL to batch and deduplicate issued queries before making the actual network calls.

// Independent computations

val resultAB = rollupAB(metricA(), metricB())
val resultAC = rollupAC(metricA(), metricC())

// Independent computations val mA = val resultAB = rollupAB(metricA(), metricB()) val resultAC = rollupAC(metricA(), metricC())

// Optimize: common subexpression elimination
val mA = metricA()
val resultAB = rollupAB(mA, metricB())
val resultAC = rollupAC(mA, metricC())

Common subexpression elimination

ŻΑ

In compiler theory, **common subexpression elimination** (CSE) is a compiler optimization that searches for instances of identical expressions (i.e., they all evaluate to the same value), and analyzes whether it is worthwhile replacing them with a single variable holding the computed value.

 \equiv Contents \checkmark

^ Example

☆

In the following code:

a = b * c + g; d = b * c * e;

it may be worth transforming the code to:

tmp = b * c; a = tmp + g; d = tmp * e;

if the cost of storing and retrieving tmp is less than the cost of calculating b * c an extra time.

Let's Write an Optimizing EDSL Compiler

Observations

Semantically-parallel Applicative instances don't bode well with Monad instances. If a type is a monadic, a lawful Applicative instance has to be sequential.

This is why Cats exposes a **Parallel** type-class, inspired by PureScript, which allows client code to choose between semantically-sequential and semantically-parallel Applicative instances.

Future Work

Future Work Static Analysis on Arrow Computations

Call for Presentations ionut.g.stan@gmail.com

Questions!

Thanks!