

The Quantum Mechanics of Data Pipelines

Pere Urbon-Bayes

Data Wrangler

pere.urbon @ { gmail.com, [acm.org](mailto:pere.urbon@acm.org) }

<http://www.purbon.com>

Who am I

Pere Urbon-Bayes (Berliner since 2011)

Software Architect and Data Engineer

Passionate about Systems, Data and Teams

Free and Open Source Advocate and Contributor

Working on a Data Engineering book for everybody

Software Architect and Data Engineer for Hire



Springer Nature in Berlin



Our Brands

 Springer

nature
research

 **BMC**

palgrave
macmillan

**SCIENTIFIC
AMERICAN**

Apress®

 **Adis**

 **J.B. METZLER**

 **macmillan**
education

 Springer Healthcare

Topics of Today

- Making data available across the company.
- From the warehouse to the era of real time.
 - Approaches to make data available.
 - Benefits and Challenges.
- The hardest problem, the human part.



Information systems, sharing
data between applications
since last century....

Making systems communicate

A black and white portrait of an elderly man with a long, full white beard and hair. He is looking slightly to the left. The background is dark. The text "A totally random system evolutionary tale" is overlaid in orange at the bottom.

**A totally random system
evolutionary tale**

L E O N A R D O D I C A P R I O
KEN WATANABE JOSEPH GORDON-LEVITT MARION COTILLARD ELLEN PAGE TOM HARDY CILLIAN MURPHY TOM BERENGER AND MICHAEL CAINE

FROM THE DIRECTOR OF THE DARK KNIGHT

INCEPTION

EXPERIENCE IT **JULY 16** IN THEATERS AND **IMAX**

PG-13 PARENTS STRONGLY CAUTIONED
Some Material May Be Inappropriate for Children Under 13
Reasons of Violence and Action Throughout

inceptionmovie.com

LEGENDARY
PICTURES

SYNCOPE

WARNER BROS. PICTURES
©2010 Warner Bros. Ent. All Rights Reserved





The Analyst Emergence

A photograph of a server room with a dense network of cables and electronic equipment. In the foreground, a thick, braided white cable runs vertically. To its left, a black cable is plugged into a network port on a white device. A label on the device reads "Environnement" and "TYPE AC32M". In the background, there are more cables, some with green status lights, and a blue metal rack. The overall scene is a complex web of wires and hardware.

The easiest system is the isolated one

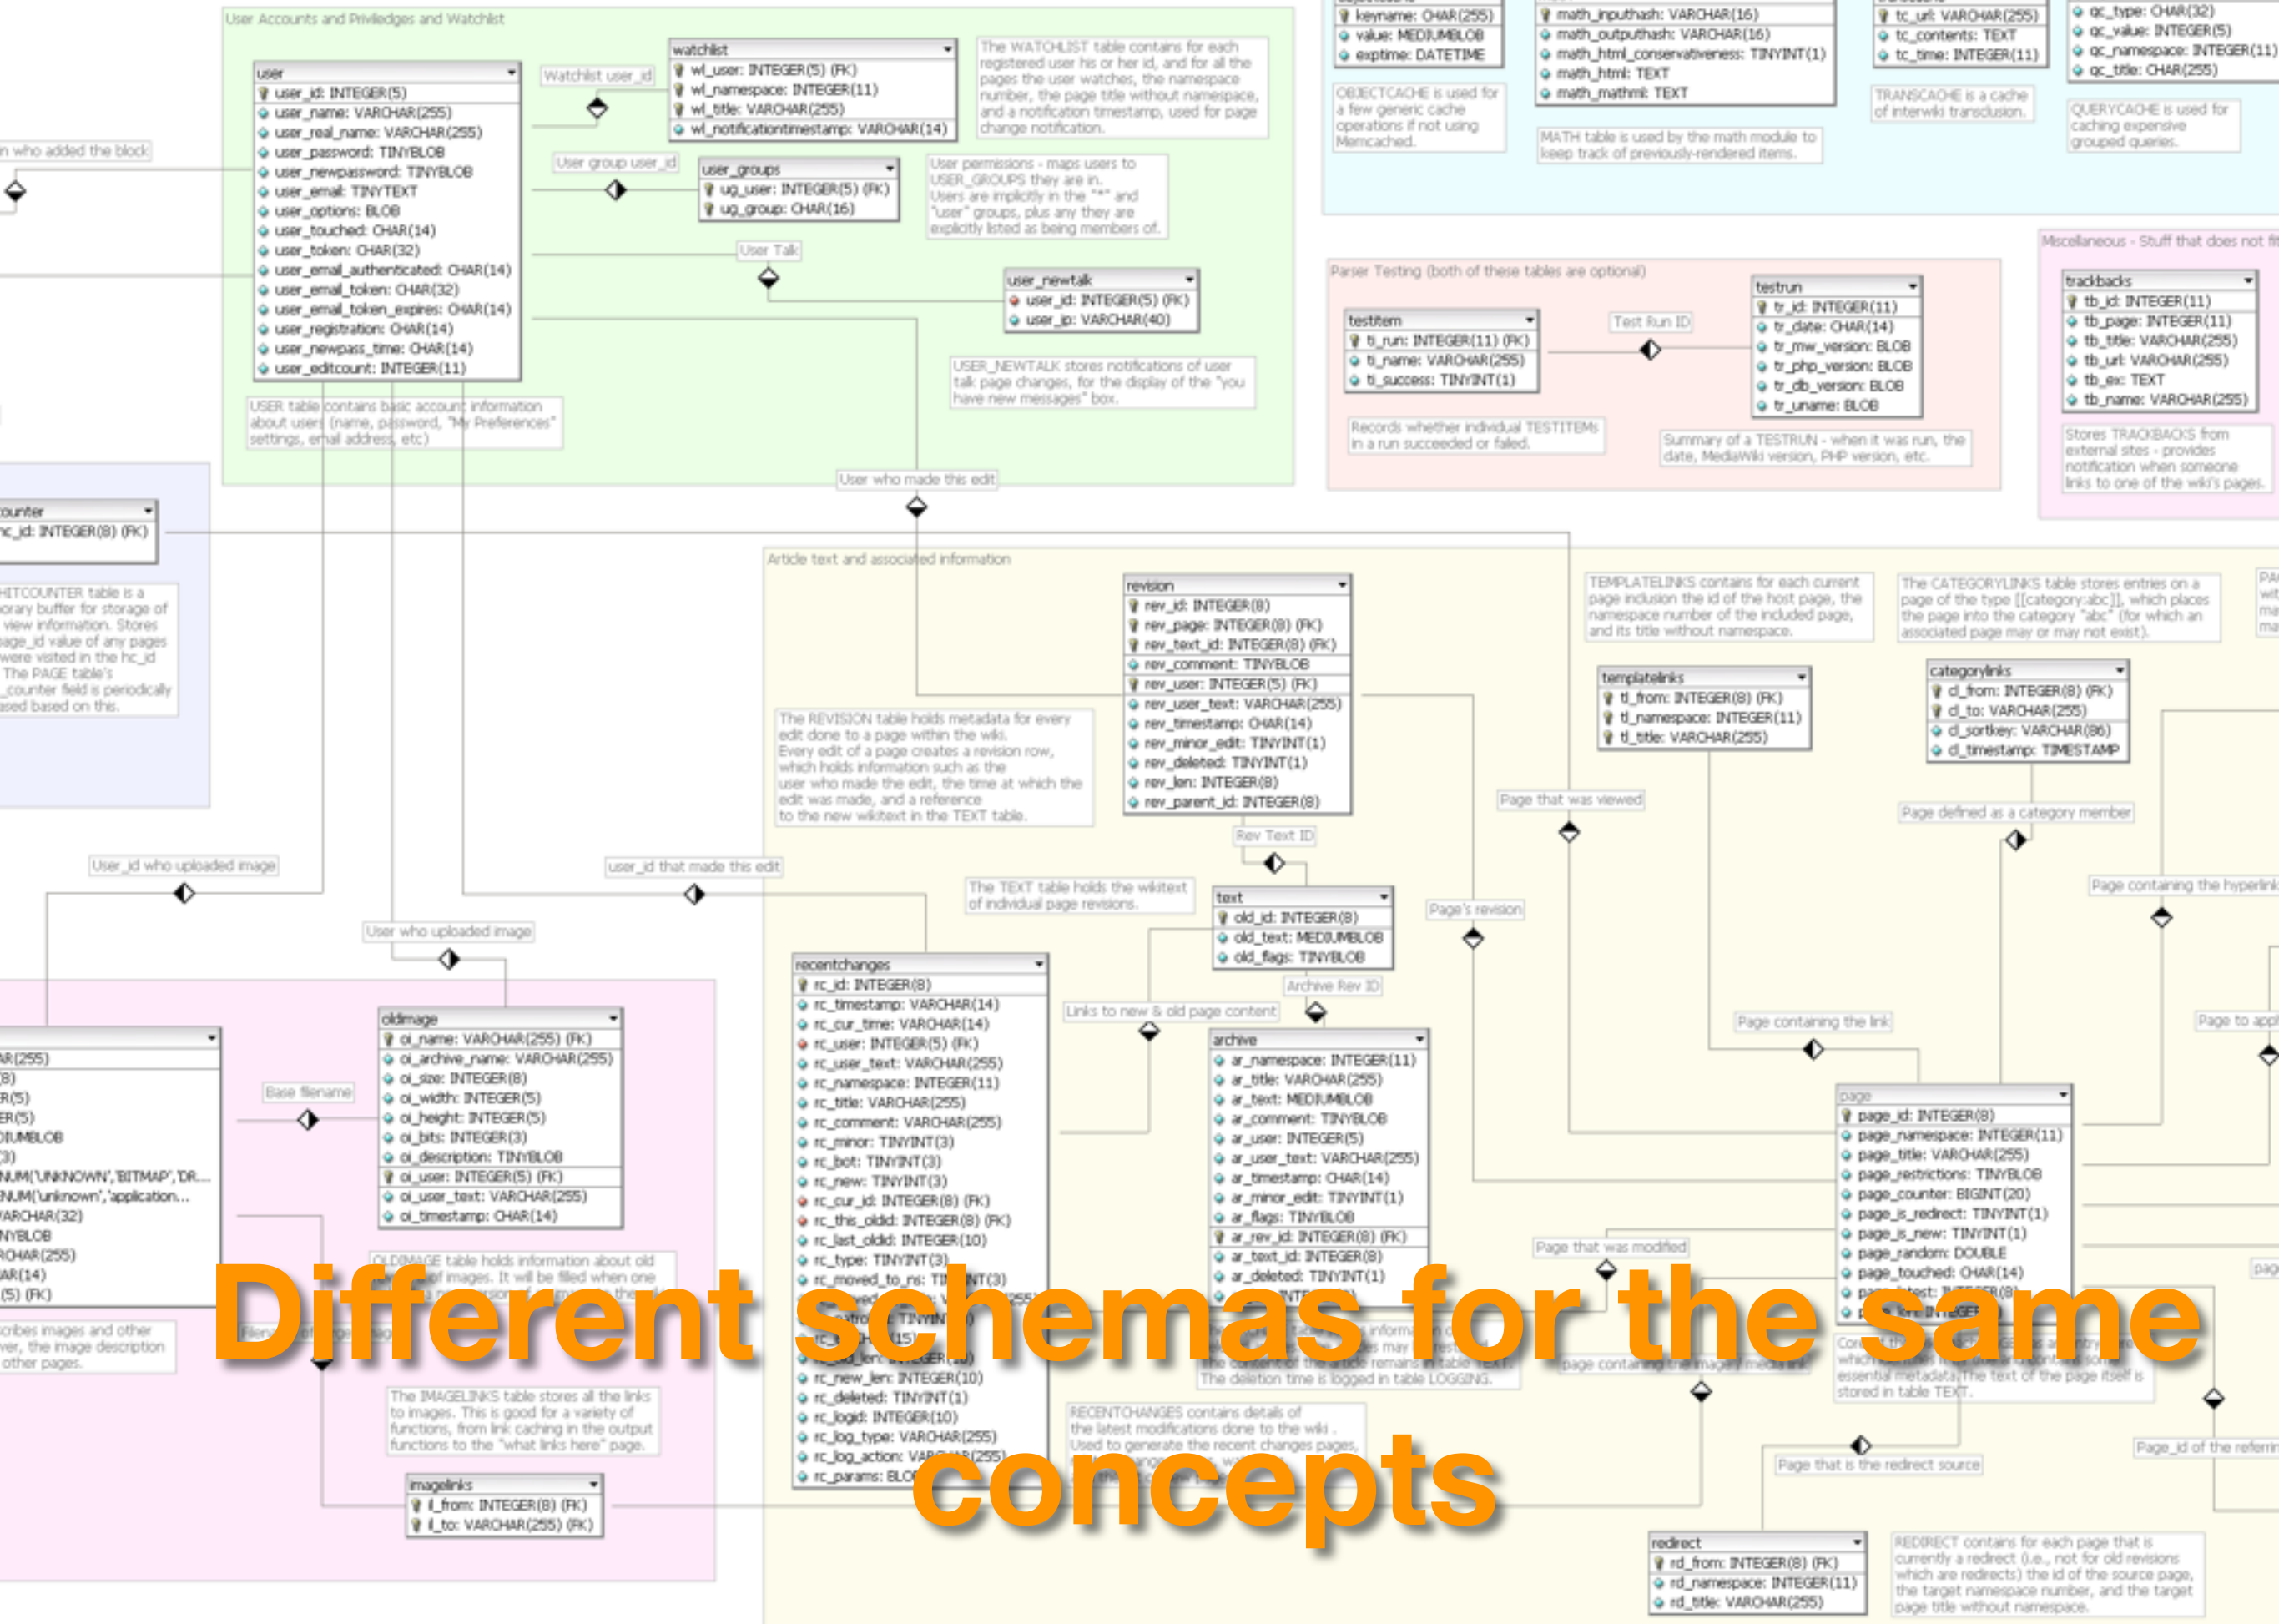


Acquiring or being acquired



**The ever growing chaos of a
technology variety**

for specific details on each of these tables.



Making data available is a system integrations problem.

The challenges in connecting data



CHANGE

Change




Dealing with failure, is hard



**The ever growing performance
battle..**





**Loosely coupled systems,
bringing maintainability to data**



Building a Shantytown, The Big Ball of Mud

"organisations which design systems ... are constrained to produce designs which are copies of the communication structures of these organisations."

– M. Conway



From the warehouse to
the real time era

Processing data at the speed of light

What is a data pipeline?



Because systems do not live in isolation anymore, they need to incorporate and/or generate data for other components.



Working with batches

The workers approach

Working in Batches

Batch processing is:

- The execution of a series of jobs/tasks
- In a computer, or group of computers.
- Without manual intervention.

A job is the single unit of work.

Working in Batches

Batches are a natural mapping from procedural and OO programming paradigms.

Implementation follow up from traditional multithread models.

An image uploaded using Sidekiq

```
class ProductImageUploader
  include Sidekiq::Worker

  def perform(image_id)
    s3_upload(data_for_image(image_id))
  end

  def self.upload(product, image_ids)
    image_ids.each do |image_id|
      perform_async(image_id)
    end
  end
end
```


Computing Pi using spark.

```
/** Computes an approximation to pi */
object SparkPi {
  def main(args: Array[String]) {
    val spark = SparkSession
      .builder
      .appName("Spark Pi")
      .getOrCreate()
    val slices = if (args.length > 0) args(0).toInt else 2
    val n = math.min(100000L * slices, Int.MaxValue).toInt // avoid overflow
    val count = spark.sparkContext.parallelize(1 until n, slices).map { i =>
      val x = random * 2 - 1
      val y = random * 2 - 1
      if (x*x + y*y <= 1) 1 else 0
    }.reduce(_ + _)
    println("Pi is roughly " + 4.0 * count / (n - 1))
    spark.stop()
  }
}
```


Common best practices

- Make your jobs **small and simple**, to ensure performance and maintainability.
- Make your jobs **idempotent and transactional**, to ensure safety and resilience.
 - At least once, exactly once,....
- Embrace **concurrency and asynchronous** api's to bring utilisation to the top.



**The pros and cons of
this approach**



Building pipelines to
transport data

Data plumbers since 1983

Working in streams

Stream processing is a computer paradigm that

- Provides a simplified parallel computation methodology.
- Given a stream of data, a series of (pipelined) operations can be applied.

Streams power algorithmic trading, RFID's, fraud detection, monitoring, telecommunications and many more.


Working in streams

Related paradigms are:

- **Data Flow:** A program as data flowing between operations.
- **Event Stream Processing:** Databases, Visualisation, middleware and languages to build event based apps.
- **Reactive Programming:** Async programming paradigm concerned with data streams and propagation of change.


Data Flow Programming

- Model programs as a DAG graph of data flowing between operations.
- Data flow through databases, brokers, streams...
- Operation types:
 - Enrichment
 - Drop/Throttle
 - Transform
- Backpressure, buffers, reactive.




Backpressure

- It describes the build-up of data behind an I/O switch if the buffers are full and incapable of receiving any more data.
- The transmitting device halts the sending of data packets until the buffers have been emptied and are once more capable of storing information.




Backpressure

- It describes the build-up of data behind an I/O switch if the buffers are full and incapable of receiving any more data.
- The transmitting device halts the sending of data packets until the buffers have been emptied and are once more capable of storing information.



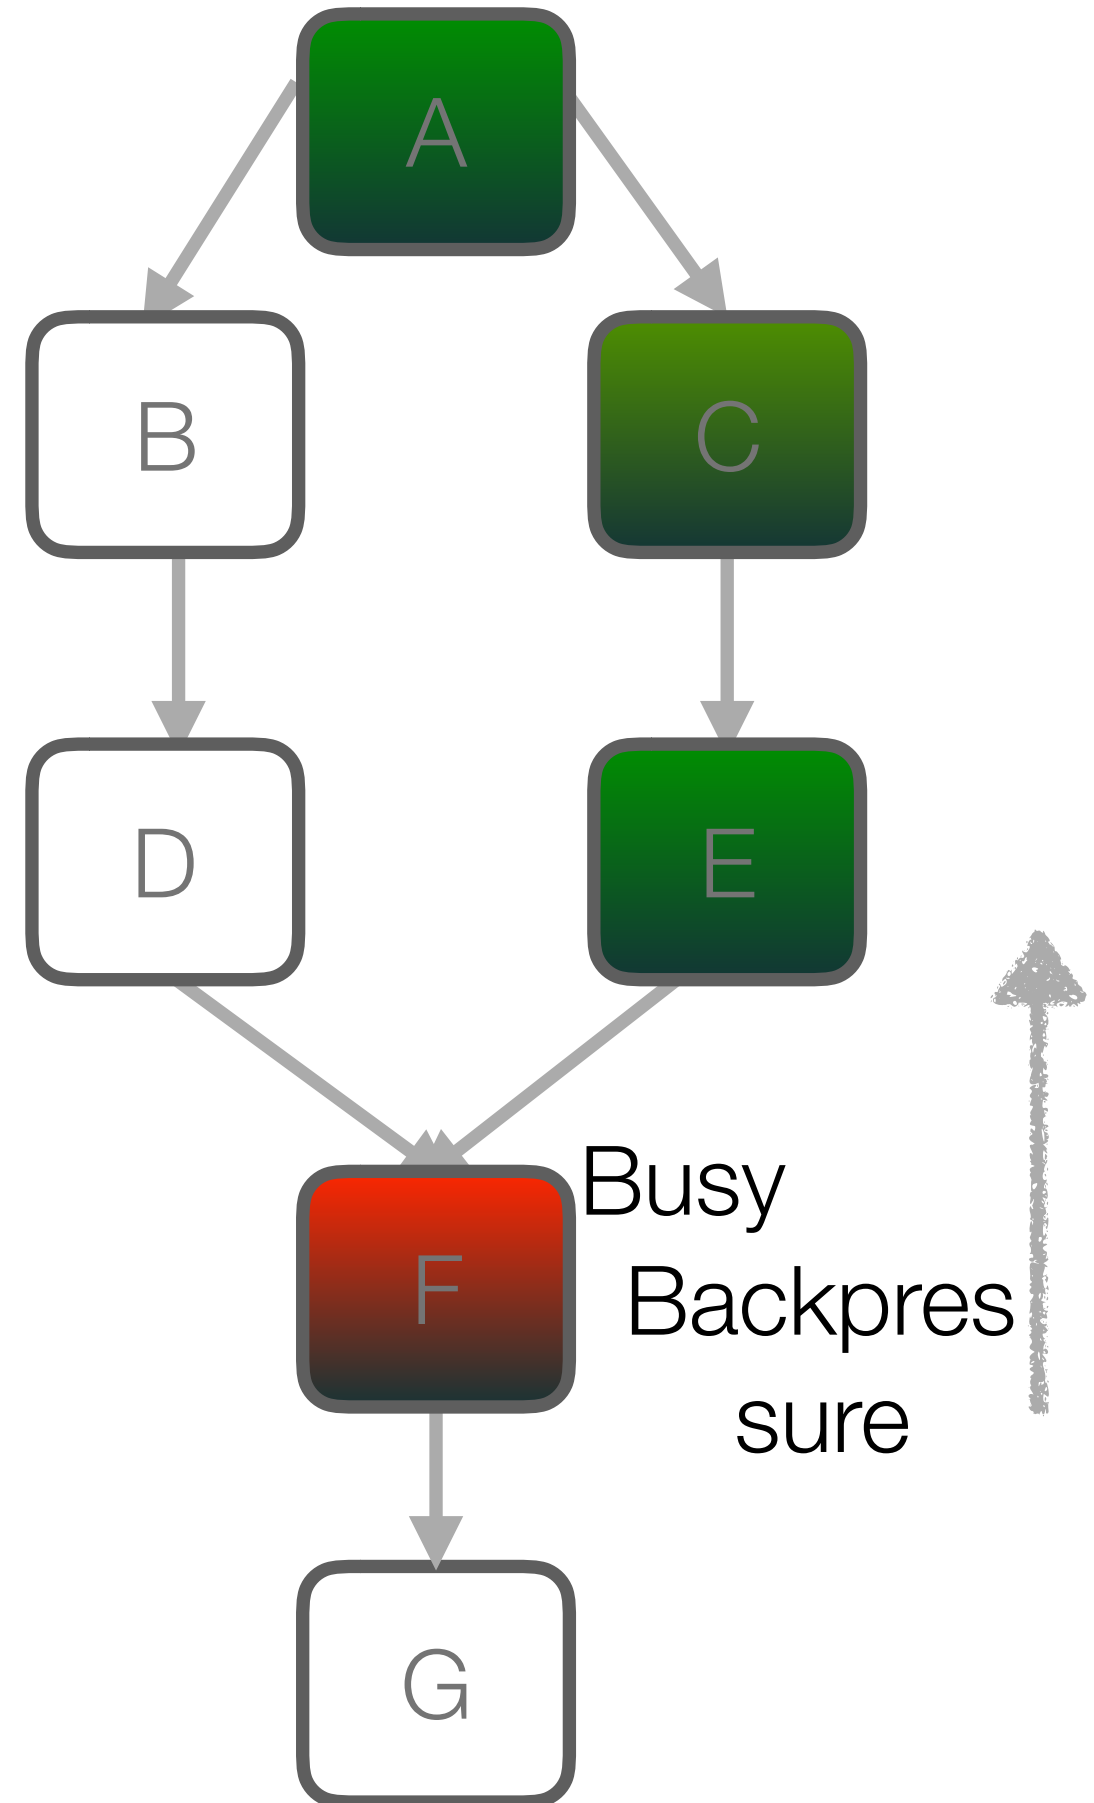
Backpressure

- It describes the build-up of data behind an I/O switch if the buffers are full and incapable of receiving any more data.
- The transmitting device halts the sending of data packets until the buffers have been emptied and are once more capable of storing information.



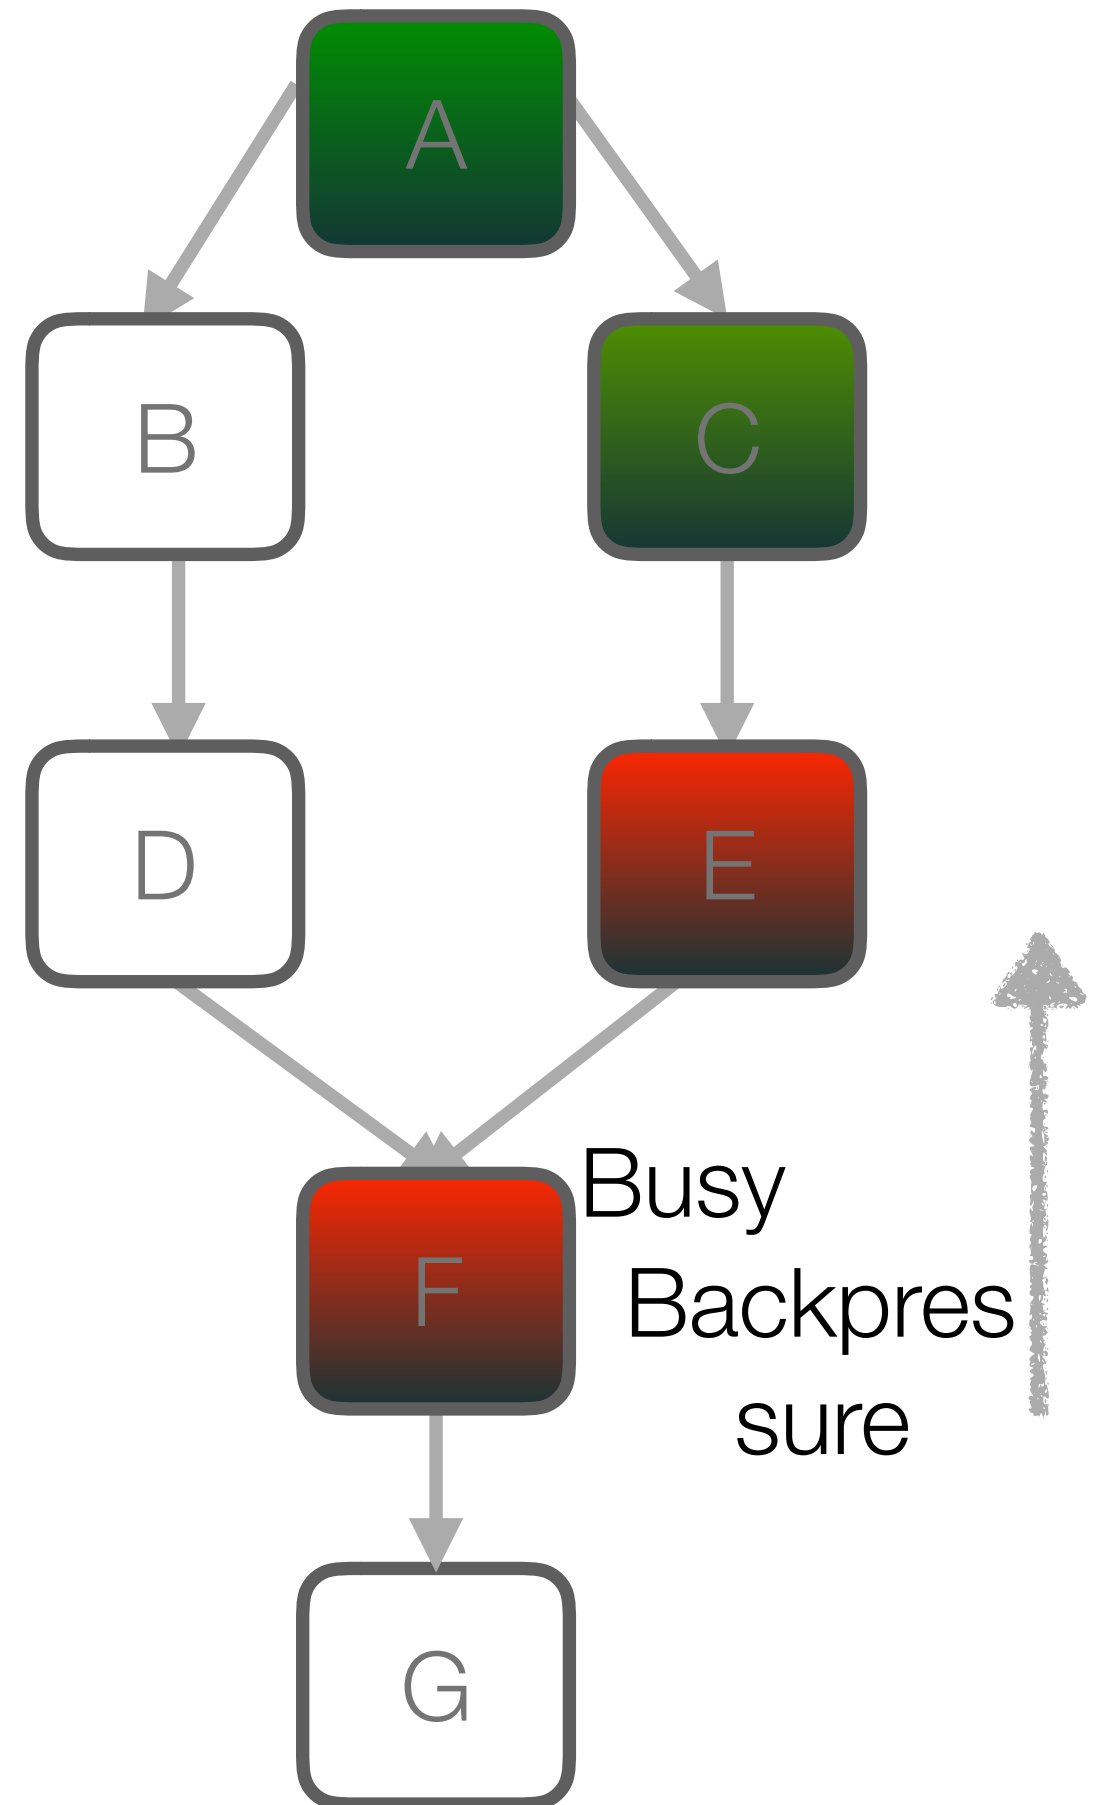
Backpressure

- It describes the build-up of data behind an I/O switch if the buffers are full and incapable of receiving any more data.
- The transmitting device halts the sending of data packets until the buffers have been emptied and are once more capable of storing information.




Backpressure

- It describes the build-up of data behind an I/O switch if the buffers are full and incapable of receiving any more data.
- The transmitting device halts the sending of data packets until the buffers have been emptied and are once more capable of storing information.



Backpressure

- It describes the build-up of data behind an I/O switch if the buffers are full and incapable of receiving any more data.
- The transmitting device halts the sending of data packets until the buffers have been emptied and are once more capable of storing information.



Reactive Streaming

When one component is struggling to keep-up, the system as a whole needs to respond in a sensible way. It is unacceptable for the component under stress to fail catastrophically or to drop messages in an uncontrolled fashion. Since it can't cope and it can't fail it should communicate the fact that it is under stress to upstream components and so get them to reduce the load.

<http://www.reactive-streams.org/>

<http://www.reactivemanifesto.org/>

A word count on reddit with Akka Streams.

```
def main(args: Array[String]): Unit = {  
  // 0) Create a Flow of String names, using either  
  //    the argument vector or the result of an API call.  
  val subreddits: Source[String, Unit] =  
    if (args.isEmpty)  
      Source(RedditAPI.popularSubreddits).mapConcat(identity)  
    else  
      Source(args.toVector)  
  
  val res: Future[Map[String, WordCount]] =  
    subreddits  
      .via(fetchLinks)  
      .via(fetchComments)  
      .runWith(wordCountSink)  
  
  res.onComplete(writeResults)  
  
  as.awaitTermination()  
}
```



A word count on reddit with Akka Streams.


```
val fetchLinks: Flow[String, Link, Unit] =  
  Flow[String]  
    .via(throttle(redditAPIRate))  
    .mapAsyncUnordered( subreddit => RedditAPI.popularLinks(subreddit) )  
    .mapConcat( listing => listing.links )  
  
val fetchComments: Flow[Link, Comment, Unit] =  
  Flow[Link]  
    .via(throttle(redditAPIRate))  
    .mapAsyncUnordered( link => RedditAPI.popularComments(link) )  
    .mapConcat( listing => listing.comments )
```



A word count on reddit with Akka Streams.


```
val wordCountSink: Sink[Comment, Future[Map[String, WordCount]]] =  
  Sink.fold(Map.empty[String, WordCount])(  
    (acc: Map[String, WordCount], c: Comment) =>  
      mergeWordCounts(acc, Map(c.subreddit -> c.toWordCount))  
  )
```


Indexing into Solr with Apache NiFi

	 GetTwitter GetTwitter 1.3.0 org.apache.nifi - nifi-social-media-nar	
In	0 (0 bytes)	5 min
Read/Write	0 bytes / 0 bytes	5 min
Out	0 (0 bytes)	5 min
Tasks/Time	0 / 00:00:00.000	5 min

	Extract lang & text EvaluateJsonPath 1.3.0 org.apache.nifi - nifi-standard-nar	
In	0 (0 bytes)	5 min
Read/Write	0 bytes / 0 bytes	5 min
Out	0 (0 bytes)	5 min
Tasks/Time	0 / 00:00:00.000	5 min

	Route Non-Empty English Tweets RouteOnAttribute 1.3.0 org.apache.nifi - nifi-standard-nar	
In	0 (0 bytes)	5 min
Read/Write	0 bytes / 0 bytes	5 min
Out	0 (0 bytes)	5 min
Tasks/Time	0 / 00:00:00.000	5 min

	MergeContent MergeContent 1.3.0 org.apache.nifi - nifi-standard-nar	
In	0 (0 bytes)	5 min
Read/Write	0 bytes / 0 bytes	5 min
Out	0 (0 bytes)	5 min
Tasks/Time	0 / 00:00:00.000	5 min

	PutSolrContentStream PutSolrContentStream 1.3.0 org.apache.nifi - nifi-solr-nar	
In	0 (0 bytes)	5 min
Read/Write	0 bytes / 0 bytes	5 min
Out	0 (0 bytes)	5 min
Tasks/Time	0 / 00:00:00.000	5 min

Name success
Queued 0 (0 bytes)

Name matched
Queued 0 (0 bytes)

Name tweet
Queued 0 (0 bytes)

Name merged
Queued 0 (0 bytes)

Name connection_failure
Queued 0 (0 bytes)

Streaming projects



Apache **Kafka**

Apache **Samza**



Challenges in data plumbing

- Systems growth is most commonly mapping internal communication channels among the organisation.
- This introduces challenges on several areas:
 - Handle failure.
 - Keep the data processes low latency.
 - Changes in communication and data.
 - Data availability and governance.



**The pros and cons of
this approach**




Scaling Human Data
Communication

Handling communication patterns

“Data pipelines emerge to automate the communications structures of the organisation”

–Your data engineer next door

Is all about communication, right?



Accessing data in a more reliable way

Problems usually pop because of changes in:

- Data expectations: When inbound teams need to change the internal data representation, volumes or schemas unexpected results are expected.
- Communication channels: A software platform is all about communication between components, also data pipelines, if they are changed users should handle it.

The shared schema registry

A centralised schema registry is a service where all organisation wide schemas are made accessible, facilitating access across teams, in detail benefits are:

- Simplify organisational data management challenges.
- Build resilient data pipelines.
- Record schema evolution.
- Facilitate data discovery across teams.
- Stream cost efficient data platforms.
- Policy enforcements.

The shared schema registry

- Popular ways to achieve this is by storing schemas in formats such as Avro, Protocol Buffers or Thirft, preferable the first one.
- Curiously there exist many private implementations, the first opensouce one is the kafka centric schema-registry by Confluent INC.
- Consumer-Driven Contracts: Similar approach introduced in 2006 by Ian Robinson at ThoughtWorks.
 - Popular implementation: pact.io

Domain Driven Design

Domain-driven design (DDD) is an approach to software development for complex needs by connecting the implementation to an evolving model.

One of the premises is the creative collaboration between technical and domain experts to refine the model.

https://en.wikipedia.org/wiki/Domain-driven_design

References

- Pat Helland. Accountants don't use erasers. <https://blogs.msdn.microsoft.com/pathelland/2007/06/14/accountants-dont-use-erasers/>
- Martin Kleppmann. Turning databases inside out. <https://www.confluent.io/blog/turning-the-database-inside-out-with-apache-samza/>
- Martin Kleppman. Designing Data Intensive Applications. <https://dataintensive.net/> .O'Reilly.

References

- Gregor Hohpe. Enterprise Integration Patterns.
<http://www.enterpriseintegrationpatterns.com/>
- Matt Welsh, et al. SEDA: An Architecture for Well-Conditioned, Scalable Internet Services.
<http://www.sosp.org/2001/papers/welsh.pdf>

LOONEY TUNES



© 1990 Warner Bros. Entertainment Inc.

Thanks a lot, Questions?

The Quantum Mechanics of Data Pipelines

Pere Urbon-Bayes

Data Wrangler

pere.urbon @ { gmail.com, [acm.org](mailto:pere.urbon@acm.org) }

<http://www.purbon.com>