Microservices with CARGO



Vikram Nitin Columbia University



Shubhi Asthana **IBM Research**

IBM Research

https://arxiv.org/abs/2207.11784



Migrating Monolithic Applications to



Baishakhi Ray Columbia University



Rahul Krishna IBM Research





GitHub's Slow March toward Monolithic Hell



2008 Github is created with a monolithic Ruby on Rails architecture



IBM Research



Why did it take that much effort do a version upgrade?



Monolithic code can be problematic!

"...having everyone doing development in the same monolithic code base is no longer the most efficient and optimal way to scale GitHub."

Sha Ma, VP of Software Engineering at **Github** "... over time, as that project matures, as you add more developers on it, as it grows ... that monolith is going to add overhead into your process, and that software development lifecycle is going to begin to slow down."

https://www.infoq.com/presentations/github-rails-monolith-microservices/
 https://thenewstack.io/led-amazon-microservices-architecture/

IBM Research



Rob Brigham, senior manager at **Amazon AWS**



Large, complex code structure is difficult for developers to comprehend





^[1] <u>Chris Richardson, "Microservice Patterns", 2018</u>

IBM Research





Large, complex code **structure** is difficult for developers to comprehend Development is slow because **build time** and start up time are high







^[1] <u>Chris Richardson, "Microservice Patterns", 2018</u>

IBM Research





Large, complex code **structure** is difficult for developers to comprehend Development is slow because **build time** and start up time are high

Deploying into production is a long process, so continuous development is difficult







^[1] <u>Chris Richardson, "Microservice Patterns", 2018</u>

IBM Research





Large, complex code **structure** is difficult for developers to comprehend Development is slow because **build time** and start up time are high



Deploying into production is a long process, so continuous development is difficult Scaling to higher workloads is difficult









^[1] <u>Chris Richardson, "Microservice Patterns", 2018</u>

IBM Research









a long process, so continuous development is difficult

Scaling to higher workloads is difficult



Locked into **obsolete** technology stack. Difficult to adopt new frameworks









IBM Research



From Monoliths to Microservices



^[1] <u>Chris Richardson, "Microservice Patterns", 2018</u>





Each microservice :



Has code with **single functionality**



^[1] <u>Chris Richardson, "Microservice Patterns", 2018</u>



Each microservice :



Has code with **single** functionality

Has a **private database** on which it performs transactions



Restaurant

^[1] <u>Chris Richardson, "Microservice Patterns", 2018</u>



Each microservice :



Has code with **single functionality**



Has a **private database** on which it performs transactions



Communicates with the others either **asynchronously** (message bus, queues, etc.) or **synchronously** (REST API, gRPC, etc.)



^[1] <u>Chris Richardson, "Microservice Patterns", 2018</u>



Each microservice :

Has code with **single** functionality

Has a **private database** on



which it performs transactions Communicates with the others either **asynchronously** (message bus, queues, etc.) or synchronously (REST API,

gRPC, etc.)

Is highly **cohesive** and loosely **coupled** with the other services









Decomposing Monoliths into Microservices

This is a hard problem!

Identifying functional boundaries requires considerable 1. domain expertise



Decomposing Monoliths into Microservices

This is a hard problem!

- Identifying functional boundaries requires considerable 1. domain expertise
- 2. Separating classes according to functionality is not sufficient:
 - May lead to **Concurrency** issues Ι.
 - ii. Maintaining data consistency is challenging



Decomposing Monoliths into Microservices

This is a hard problem!

- Identifying functional boundaries requires considerable 1. domain expertise
- 2. Separating classes according to functionality is not sufficient:
 - May lead to **Concurrency** issues Ι.
 - Maintaining data consistency is challenging ii.



Automated monolith decomposition tools



Automated Monolith Decomposition

Active area of Research

| 1 | Escobar, D. et al.: Towards the understanding and evolution of monolithic applications as microservices In: Proceedings of 42nd Latin American Computing Conference, CLEI. (2016) |
|---|--|
| 2 | Levcovitz, A. et al.: Towards a Technique for Extracting Microservices from Monolithic Enterprise Systems. In: 3rd Brazilian Workshop on Software Visualization, Evolution and Maintenance (VEM) .pp. 97–104 (2015) |
| 3 | Ahmadvand, M., Ibrahim, A.: Requirements reconciliation for scalable and secure microservice (de)composition. In: Proceedings - 2016 IEEE 24th International Requirements Engineering Conference Workshops, REW 2016. pp. 68–73 (2016) |
| 4 | Baresi, L. et al.: Microservices Identification Through Interface Analysis. In: ESOCC 2017: Service- Oriented and Cloud Computing. pp. 19–33 (2017) |
| 5 | Gysel, M. et al.: Service cutter: A systematic approach to service decomposition. In: Lecture Notes in Computer Science. pp. 185–200 (2016) |
| 6 | Mazlami, G. et al.: Extraction of Microservices from Monolithic Software Architectures. In: 2017 IEEE International Conference on Web Services (ICWS). pp. 524–531 (2017) |
| 7 | Mustafa, O., Gómez, J.M.: Optimizing economics of microservices by planning for granularity level Experience Report. (2017) |
| 8 | Hassan, S. et al.: Microservice Ambients: An Architectural Meta-Modelling Approach for |

[1] Fritzsch, Jonas, et al. "From monolith to microservices: A classification of refactoring approaches."

IBM Research

Industrial tools

[1] https://www.ibm.com/cloud/mono2micro

vFunction

[2] <u>https://www.konveyor.io</u>

KONVEYOR

[3] <u>https://vfunction.com</u>



Automated monolith decomposition tools

Monolithic Application





Automated monolith decomposition tools







Automated monolith decomposition tools







IBM Research

A real example from a benchmark application, Daytrader. This is a portion of a call-graph, and there is a call edge between the classes TradeDirect and TradeConfig.



In this commonly recommended partitioning, **QuoteDatabean** and **TradeDirect** lie in different partitions.





But if we look beyond the call graph, we find this...





Some challenges in implementing this partitioning scheme :

Distributed monolith : The two classes 1. **QuoteDatabean** and **TradeConfig** are tightly coupled (access shared heap objects)





Some challenges in implementing this partitioning scheme :

- 1. **Distributed monolith :** The two classes **QuoteDatabean** and **TradeConfig** are tightly coupled (access shared heap objects)
- 2. Distributed transaction : QuoteDatabean and **TradeConfig** write to the same DB.





In reality, **QuoteDatabean** and **TradeDirect** are tightly coupled! Our algorithm, CARGO, groups them in the same partition





How do we get better partitions?

Our algorithm CARGO, uses the following key ideas :

classes and build a program dependency graph (PDG)

- More complete static analysis: Capture many types of dependencies between



How do we get better partitions?

Our algorithm CARGO, uses the following key ideas :

classes and build a program dependency graph (PDG) 2

PDG

- More complete static analysis: Capture many types of dependencies between
- **Explicitly model transactions:** Database transactions are added as edges in the





How do we get better partitions?

Our algorithm CARGO, uses the following key ideas:

classes and build a program dependency graph (PDG)

Explicitly model transactions: Database transactions are added as edges in the 2 PDG.

3 algorithm to assign partitions to nodes in the PDG.

IBM Research

- More complete static analysis: Capture many types of dependencies between

Detect communities in the PDG: We present a novel community detection







Our System - Step 1

Monolith



Extract PDG









| 1 2 3 4 | void main(0 A a1 = ne Object v2 |
|----------------------------|---|
| 5 6 7 | A a2 = ne Object v2 } |
| | |
| 1 2 3 4 5 6 | <pre>class A { Object fo B b = 0 return } }</pre> |
| | |
| 1 2 3 4 5 6 | <pre>class B { Object back } }</pre> |

```
(Object[] args) {
new A();
/1 = a1.foo(new Object());
new A();
/2 = a2.foo(new Object());
```

```
boo(Object v) {
  new B();
  b.bar(v);
```

```
par(Object v) {
```



Context-insensitive Analysis

main()

| 1 2 3 4 5 6 7 | <pre>void main(Object[] a: A a1 = new A(); // Object v1 = a1.foo A a2 = new A(); // Object v2 = a2.foo }</pre> |
|---------------------------------|--|
| | |
| | |
| 1 | class A { |
| 2 | Object too(Object |
| 3 | B b = new B(); |
| 4 | return b.bar(v); |
| 5 | 5 |
| 6 | 5 |
| | |
| | |
| 1 | class B { |
| 2 | Ubject par(Ubject |
| З л | ζ. |
| 4 | ۲ ۲ |
| 6 | |
| | |

IBM Research

```
Object[] args) {
ew A(); // A/1
l = a1.foo(new Object());
ew A(); // A/2
2 = a2.foo(new Object());
po(Object v) {
```

Context-sensitive Analysis

[Φ, Φ]

main()

ar(Object v) {





Context-insensitive Analysis



| 1 2 3 4 5 6 7 | <pre>void main(0 A a1 = ne Object v1 A a2 = ne Object v2 }</pre> |
|---------------------------------|--|
| 1 2 3 4 5 6 | <pre>class A { Object fo B b = r return } }</pre> |
| 1 2 3 4 5 6 | <pre>class B { Object ba } }</pre> |

IBM Research



Context-sensitive Analysis







Context-insensitive Analysis



| 1 2 3 4 5 6 7 | <pre>void main(0 A a1 = ne Object v1 A a2 = ne Object v2 }</pre> |
|---------------------------------|--|
| 1 2 3 4 5 6 | <pre>class A { Object fo B b = r return } }</pre> |
| 1 2 3 4 5 6 | <pre>class B { Object ba } }</pre> |









Context-*insensitive* Analysis



| 1 2 3 4 5 6 7 | <pre>void main(0 A a1 = ne Object v1 A a2 = ne Object v2 }</pre> |
|---------------------------------|--|
| 1 2 3 4 5 6 | <pre>class A { Object fo B b = r return } }</pre> |
| 1 2 3 4 5 6 | <pre>class B { Object ba } }</pre> |







Building a Program Dependency Graph








IBM Research



Call-return dependency





























Our System - Step 2



IBM Research

Context Snapshots



Context-sensitive Analysis


```
1 void main(Object[] args) {
2     A a1 = new A();
3     Object v1 = a1.foo(new Object());
4 
5     A a2 = new A();
6     Object v2 = a2.foo(new Object());
7  }
```

```
1 class A {
2     Object foo(Object v) {
3         B b = new B();
4         return b.bar(v);
5     }
6     }
```





Context Snapshots

The context-sensitive PDG is a **superposition** of all possible contexts.

At any time, A.foo() and B.bar() can exist in *any one context but not both simultaneously*.





Context Snapshots



IBM Research



()



Transactional Snapshot



IBM Research

46

46



Our System - Step 3



IBM Research

Context-sensitive Label Propagation



Initial State

Magenta and Blue are the two categories of label



IBM Research

We start with some **initial assignment** of labels to nodes



Initial State

Magenta and Blue are the two categories of label



IBM Research

We start with some **initial assignment** of labels to nodes









IBM Research

Each node is assigned the majority label of its neighbors









IBM Research

Repeat until convergence (no more node updates)





Label Propagation on Transactional Snapshot

IBM Research

Label Propagation on Context Snapshots





Label Propagation on Transactional Snapshot

IBM Research

Label Propagation on Context Snapshots



Initialize labels



IBM Research

54



Transactional Snapshot





Propagate Labels







Context Snapshot 1





Propagate Labels





Context Snapshot 2





Propagate Labels





And so on for all context snapshots







Our System - Overview



IBM Research

62



Our System - Overview







Our System - Overview











| Application | Description | Java Framework | # Classes | # SQL |
|--------------|---------------------------------|-----------------------------------|-----------|-------|
| Daytrader | Trading application | Java EE 8, Websphere | 109 | 6 |
| Plants | Online plant shopping | Java EE 7, Websphere | 33 | • |
| AcmeAir | Website of a fictitious airline | Openliberty, Websphere eXtreme | 66 | ٠ |
| JPetStore | Online pet supply store | Spring, Springboot | 37 | ٩ |
| Proprietary1 | Proprietary app | • | 82 | • |



| Application | Description | Java Framework | # Classes | # SQL |
|--------------|---------------------------------|-----------------------------------|-----------|-------|
| Daytrader | Trading application | Java EE 8, Websphere | 109 | 6 |
| Plants | Online plant shopping | Java EE 7, Websphere | 33 | ● |
| AcmeAir | Website of a fictitious airline | Openliberty, Websphere eXtreme | 66 | • |
| JPetStore | Online pet supply store | Spring, Springboot | 37 | ٩ |
| Proprietary1 | Proprietary app | | 82 | • |



| Application | Description | Java Framework | # Classes | # SQL |
|--------------|---------------------------------|-----------------------------------|-----------|-------|
| Daytrader | Trading application | Java EE 8, Websphere | 109 | 6 |
| Plants | Online plant shopping | Java EE 7, Websphere | 33 | • |
| AcmeAir | Website of a fictitious airline | Openliberty, Websphere eXtreme | 66 | • |
| JPetStore | Online pet supply store | Spring, Springboot | 37 | • |
| Proprietary1 | Proprietary app | • | 82 | • |



| Application | Description | Java Framework | # Classes | # SQL |
|--------------|---------------------------------|-----------------------------------|-----------|-------|
| Daytrader | Trading application | Java EE 8, Websphere | 109 | 6 |
| Plants | Online plant shopping | Java EE 7, Websphere | 33 | ٠ |
| AcmeAir | Website of a fictitious airline | Openliberty, Websphere eXtreme | 66 | ٠ |
| JPetStore | Online pet supply store | Spring, Springboot | 37 | ٠ |
| Proprietary1 | Proprietary app | • | 82 | • |



Evaluation Setup - Baseline Approaches

| Approach | |
|-------------|--|
| Mono2Micro* | Dynamic call trace |
| CoGCN | A Graph Neural N |
| FoSCI | Genetic Search-b |
| MEM | A Minimum-Spann Edit-history and se |

IBM Research

Summary

es and hierarchical clustering.

etwork and K-Means on a static call graph

based algorithm on dynamic execution traces

ning Tree based Clustering Algorithm on a graph. emantics are used to define coupling

* Enterprise scale decomposition tool



Evaluation Setup - "refining" partitions

CARGO can be used to **refine** the partitions produced by other approaches.

O Denoted by "++" suffix.

E.g., **Mono2Micro++** denotes running CARGO with initial partition labels produced by Mono2Micro.

| Original Approach | Refined with CARGO |
|-------------------|---------------------------|
| Mono2Micro | Mono2Micro++ |
| CoGCN | CoGCN++ |
| FoSCI | FoSCI++ |
| MEM | MEM++ |



Research Questions

RQ-1 Effectiveness in remediating distributed transactions

RQ-2 Latency and Throughput improvements resulting from refined microservice partitions

RQ-3 Quality of microservice partition architectural metrics



Research Questions

RQ-1 Effectiveness in remediating distributed transactions

RQ-2 Latency and Throughput improvements resulting from refined microservice partitions

RQ-3 Quality of microservice partition architectural metrics


To minimize distributed transactions, we would like, *to the extent possible*, for each database table to be accessed from one microservice partition only

73



To minimize distributed transactions, we would like, to the extent possible, for each database table to be accessed from one microservice partition only

Evaluation

• Use Transactional Purity (TXP) to measure the tendency of a database table to be accessed by *multiple* microservices



To minimize distributed transactions, we would like, to the extent possible, for each database table to be accessed from one microservice partition only

Evaluation

- Use Transactional Purity (TXP) to measure the tendency of a database table to be accessed by *multiple* microservices
 - TXP = 1 -

$$\sum_{i=0}^{K} p_i \cdot log \left[\frac{1}{p_i}\right]$$



To minimize distributed transactions, we would like, to the extent possible, for each database table to be accessed from one microservice partition only

Evaluation

Use Transactional Purity (TXP) to measure the tendency of a database table to be accessed by <u>multiple</u> microservices

TXP = 1 -

Lower purity indicates accesses from more microservices
Higher purity indicates accesses from less microservices

$$\sum_{i=0}^{K} p_i \cdot log \left[\frac{1}{p_i}\right]$$



Daytrader



Summary

- For each of the 4 baselines, the refined partitions have higher transactional purity.
- FoSCI++ and MEM++ have transactional purity of 1.0, (i.e., no distributed transactions after repartitioning).
- CARGO (unsupervised) natively achieves transactional purity of 1.0

++ implies refinement with CARGO

CARGO 77





Research Questions

RQ-1 Effectiveness in remediating distributed transactions

RQ-2 Latency and Throughput improvements resulting from refined microservice partitions

RQ-3 Quality of microservice partition architectural metrics



Minimizing distributed transactions can offer significant runtime benefits in terms of reduced latency and improved throughput.

Evaluation

• Deploy two variants of the applications:

- 1. Application with original partitioning (Mono2Micro)¹
- 2. Application with partitions refined with CARGO (aka. Mono2Micro++)²



¹ <u>https://github.com/vikramnitin9/tackle-data-gravity-insights/tree/main/RQ2/daytrader_apps/daytrader_cargo</u> ² https://github.com/vikramnitin9/tackle-data-gravity-insights/tree/main/RQ2/daytrader apps/daytrader mono2micro

terms of reduced latency and improved throughput.

Evaluation

• Deploy two variants of the applications:

- 1. Application with original partitioning (Mono2Micro)¹
- 2. Application with partitions refined with CARGO (aka. Mono2Micro++)²

• Compare two runtime performance metrics: Latency: Time between reception and completion of a request (milliseconds)

IBM Research

Minimizing distributed transactions can offer significant runtime benefits in

- Throughput: Number of successful requests honored per unit time (requests/second)



¹ <u>https://github.com/vikramnitin9/tackle-data-gravity-insights/tree/main/RQ2/daytrader_apps/daytrader_cargo</u>

² https://github.com/vikramnitin9/tackle-data-gravity-insights/tree/main/RQ2/daytrader apps/daytrader mono2micro





IBM Research

81





Significant improvements in latency and throughput. Repartitioned application has **11% lower latency** and **120% higher throughput** on average across use cases compare to the original applications.





Research Questions

RQ-1 Effectiveness in remediating distributed transactions

RQ-2 Latency and Throughput improvements resulting from refined microservice partitions

RQ-3 Quality of microservice partition architectural metrics



RQ-3 Partitions and their Architectural Quality

| METRIC | | |
|----------|----------|----------|
| Coupling | ∇ | Average |
| Cohesion | Δ | Average |
| BCP | ∇ | Purity o |
| ICP | ∇ | Inter-pa |

 ∇ Lower is better \triangle Higher is better

IBM Research

DESCRIPTION

- Coupling among partitions
- cohesion within a partition
- of Business use cases per partition.
- artition call volume





RQ-3 Partitions and their Architectural Quality **Cohesion** \triangle

Coupling ∇

| | Mono2Micro | Mono2Micro++ | CARGO | | CoGCN | CoGCN++ | CARC |
|--------------|------------|--------------|-------|--------------|-------|---------|------|
| DAYTRADER | 0.78 | 0.02 | 0.01 | DAYTRADER | 0.37 | 0.61 | |
| PLANTS | 0.31 | 0.04 | 0.05 | PLANTS | 0.39 | 0.46 | |
| ACMEAIR | 0.58 | 0.04 | 0.03 | ACMEAIR | 0.21 | 0.32 | |
| JPETSTORE | 0.77 | 0.03 | 0.03 | JPETSTORE | 0.20 | 0.24 | |
| PROPRIETARY | 0.42 | 0.03 | 0.04 | PROPRIETARY | 0.69 | 0.73 | С |
| WIN/TIE/LOSS | 5, | /0/0 | | WIN/TIE/LOSS | 5, | /0/0 | |

CARGO improves the partitioning quality (reduced coupling and increased cohesion) of other approaches and works equally well in unsupervised mode.

IBM Research

*Mono2Micro++ performs slightly better than CARGO



RQ-3 Partitions and their Architectural Quality

| | Mono2Micro | Mono2Micro++ | CARGO |
|--------------|------------|--------------|-------|
| DAYTRADER | 2.31 | 2.57 | 1.31 |
| PLANTS | 1.68 | 2.20 | 1.79 |
| ACMEAIR | 1.29 | 1.48 | 1.75 |
| JPETSTORE | 2.25 | 2.35 | 2.87 |
| PROPRIETARY | 1.53 | 1.23 | 1.55 |
| WIN/TIE/LOSS | 0/ | | |

CARGO performs poorly on BCP. The definition of BCP depends heavily on the quality of the generated business use cases, which Mono2Micro has access to but we do not.

IBM Research

BCP ∇





Partitioning monolithic code can be challenging







Existing automated approaches miss key code and/or transactional dependencies









Existing automated approaches miss key code and/or transactional dependencies



We present CARGO, which uses (a) precise static analysis, (b) explicit modeling of database transactions, and (c) a novel community detection algorithm







Partitioning monolithic code can be challenging



Existing automated approaches miss key code and/or transactional dependencies



We present CARGO, which uses (a) precise static analysis, (b) explicit modeling of database transactions, and (c) a novel community detection algorithm



Compared to existing approaches, CARGO (a) reduces distributed transactions, (b) achieves better latency and throughput, (c) achieves better performance on architectural metrics







Future Directions

- distributed patterns like **SAGA**.
- partitions.
 - vs async
 - Automatically generate IPC code

IBM Research

• Sometimes, it is not possible to avoid distributed transactions. In that case, we need mechanism to automatically refactor database transactions into

Once we identify partitions, there is still a lot of work to implement these

Identify best mechanism for Inter-Process Communication (IPC), e.g., sync





Thank You!

https://arxiv.org/abs/2207.11784

