



Type-Safe Dynamic Placement with First-Class Placed Values


George Zakhour




Pascal Weisenburger



Guido Salvaneschi






Dynamic Placement

Dynamic Placement

- Caching
- Geo-Replication
- Hybrid Cloud
- High-Performance Computing





Safe Placement

Safe Placement and Static Checking

Type-Safe Distributed Programming with ML5*

Tom Murphy VII, Karl Crary, and Robert Harper

Department of Computer Science
Carnegie Mellon University
Pittsburgh, PA, USA
{tom7, crary, rwh}@cs.cmu.edu

A Formal Theory of Choreographic Programming

Luís Cruz-Filipe¹ · Fabrizio Montesi¹ · Marco Peressotti¹

Received: 7 September 2022 / Accepted: 9 April 2023 / Published online: 27 May 2023
© The Author(s) 2023

Abstract

Choreographic programming is a paradigm for writing coordination plans for distributed systems from a global point of view, from which correct-by-construction decentralised implementations can be generated automatically. Theory of choreographies typically includes a number of complex results that are proved by structural induction. The high number of cases and the subtle details in some of these proofs has led to important errors being found in published works. In this work, we formalise the theory of a choreographic programming

language in Coq. Our development includes its Turing completeness, a compilation principle, characterisation of the correctness of this principle, the benefits of using a theorem prover: we mechanise the proof, and a significant step offer a foundation for the future formal de-

language for spa-
of ML, allows an
soned about as a
bility of any kind
gic, statically ex-
the ML5 compiler
ved in the com-
different resources
iler and runtime
ng: a distributed
web server.

Distributed System Development with SCALALoCI

PASCAL WEISENBURGER, Technische Universität Darmstadt, Germany
MIRKO KÖHLER, Technische Universität Darmstadt, Germany
GUIDO SALVANESCHI, Technische Universität Darmstadt, Germany

Distributed applications are traditionally developed as separate modules, often in different languages, which react to events, like user input, and in turn produce new events for the other modules. Separation into components requires time-consuming integration. Manual implementation of communication forces programmers to deal with low-level details. The combination of the two results in obscure distributed data flows scattered among multiple modules, hindering reasoning about the system as a whole.

The SCALALoCI distributed programming language addresses these issues with a coherent model based on placement types that enables reasoning about distributed data flows, supporting multiple software architectures via dedicated language features and abstracting over low-level communication details and data conversions. As we show, SCALALoCI simplifies developing distributed systems, reduces error-prone communication code and favors early detection of bugs.

ArchJava: Connecting Software Architecture to Implementation

Jonathan Aldrich Craig Chambers David Notkin

Department of Computer Science and Engineering
University of Washington
Box 352350
Seattle, WA 98195-2350 USA
+1 206 616-1846

{jonal, chambers, notkin}@cs.washington.edu

Abstract

Software architecture describes the structure of a system, enabling more effective design, program understanding, and formal

multiple implementation
stencils, causing
s, and inhibiting
sion to Java that
h implementation,
s to architectural
to a circuit-design
press architectural
and that it can ad-
m.

language. Thus, it may be difficult to trace architectural features to the implementation, and the implementation may become inconsistent with the architecture as the program evolves. In summary, while architectural analysis in existing ADLs may reveal important architectural properties, these properties are not guaranteed to hold in the implementation.

In order to enable architectural reasoning about an

Ur/Web: A Simple Model for Programming the Web

Adam Chlipala
MIT CSAIL
adamc@csail.mit.edu

Abstract

The World Wide Web has evolved gradually from a document delivery platform to an architecture for distributed programming. This largely unplanned evolution is apparent in the set of interconnected languages and protocols that any Web application must manage. This paper presents Ur/Web, a domain-specific, statically typed functional programming language with a much simpler model for programming modern Web applications. Ur/Web's model is unified, where programs in a single programming language are compiled to other "Web standards" languages as needed; supports novel kinds of encapsulation of Web-specific state; and exposes simple concurrency, where programmers can reason about distributed, multithreaded applications via a mix of transactions and cooperative preemption. We give a tutorial introduction to the main features of Ur/Web and discuss the language implementation and the production Web applications that use it.

Links: Web Programming Without Tiers*

Ezra Cooper, Sam Lindley, Philip Wadler, and Jeremy Yallop

University of Edinburgh

Fabric: A Platform for Secure Distributed Computation and Storage

Jed Liu Michael D. George K. Vikram
Xin Qi Lucas Wayne Andrew C. Myers
{liujed,mdgeorge,kvikram,qixin,lrw,andru}@cs.cornell.edu
Department of Computer Science
Cornell University
4130 Upson Hall, Ithaca NY

Abstract

Fabric is a new system and language for building secure distributed information systems. It is a decentralized system that allows heterogeneous network nodes to securely share both information and computation resources despite mutual distrust. Its high-level programming language makes distribution and persistence largely transparent to programmers. Fabric supports data-shipping and functional programming styles of computation: both computation and information can move between nodes to meet security requirements or to improve performance. Fabric provides a rich, Java-like object model, but data resources are labeled with confidentiality and integrity poli-

does not offer general, principled techniques for implementing the functionality of these systems while also satisfying their security and privacy requirements. This lack motivates the creation of Fabric, a platform for building secure distributed information systems.

It is particularly difficult to build secure federated systems, which integrate information and computation from independent administrative domains—each domain has policies for security and privacy, but does not fully trust other domains to enforce them. Integrating information from different domains is important because it enables new services and capabilities.

To illustrate the challenges, consider the scenario of two medical institutions that want to securely and quickly share patient information. This goal is important: according to a 1999 Institute of Medicine study, at least 44,000 deaths annually result from medical errors, with incomplete patient information identified as a leading cause [25]. However, automated sharing of patient data poses difficulties. First, the security and privacy policies of the two institutions must be satisfied (as mandated by HIPAA [22] in the U.S.), restricting which information can be shared or modified by the two institu-



time and
sure con-
gination
sults from
lean, con-
enforces

for network communication, and on a language or API like SQL for storing persistent, structured data on servers. Code fragments in these different languages are often embedded within each other in complex ways, and the popular Web development tools provide little help in catching inconsistencies.

These complaints are not new, nor are language-based solutions. The Links project [7, 11] pioneered the "tierless programming" approach, combining all the pieces of dynamic Web applications within one statically typed functional language. More recent designs in the mainstream reap some similar benefits, as in Google's Web Toolkit [3] and Closure [5] systems, for adding compilation on top of Web-standard languages; and Microsoft's LINQ [27], for type-safe querying (to SQL databases and more) within general-purpose languages.

Such established systems provide substantial benefits to Web programmers, but there is more we could ask for. This paper focuses on a language design that advances the state of the art by ad-

Safe Placement and Static Checking

Type-Safe Distributed Programming with ML5*

Tom Murphy VII, Karl Crary, and Robert Harper

Department of Computer Science
Carnegie Mellon University
Pittsburgh, PA, USA
{tom7, crary, rwh}@cs.cmu.edu

Links: Web Programming Without Tiers*

Ezra Cooper, Sam Lindley, Philip Wadler, and Jeremy Yallop

University of Edinburgh

A Formal Theory of Choreographic Programming

Luís Cruz-Filipe¹ · Fabrizio Montesi¹ · Marco Peressotti¹

Received: 7 September 2022 / Accepted: 9 April 2023 / Published online: 27 April 2023
© The Author(s) 2023

Abstract

Choreographic programming is a paradigm for writing coordinated systems from a global point of view, from which correct-by-construction implementations can be generated automatically. Theory of choreography has produced a number of complex results that are proved by structural induction and the subtle details in some of these proofs has led to important published works. In this work, we formalise the theory of choreography in Coq. Our development includes: a formalisation of its Turing completeness, a compilation principle, a characterisation of the correctness of this principle, the benefits of using a theorem prover: we mechanise the proof, and a significant step towards offering a foundation for the future formal development of choreography.

Distributed System Development with ScalaLoc

PASCAL W. SENNBURGER, Technische Universität Darmstadt, Germany
MIRKO KÖHLER, Technische Universität Darmstadt, Germany
GUIDO SALVANESCHI, Technische Universität Darmstadt, Germany

Distributed applications are traditionally developed as separate modules, often in different languages, which react to events, like user input, and in turn produce new events for the other modules. Separation into components requires time-consuming integration. Manual implementation of communication forces programmers to deal with low-level details. The combination of the two results in obscure distributed data flows scattered among multiple modules, hindering reasoning about the system as a whole.

The SCALALOCI distributed programming language addresses these issues with a coherent model based on placement types that enables reasoning about distributed data flows, supporting multiple software architectures via dedicated language features and abstracting over low-level communication details and data conversions. As we show, SCALALOCI simplifies developing distributed systems, reduces error-prone communication code and favors early detection of bugs.

ArchJava:

Connecting Software Architecture to Implementation

Jonathan Aldrich, Craig Chambers, David Nofkin

Department of Computer Science and Engineering

{jonah.aldrich, craig.chambers, nofkin}@cs.washington.edu

- Placement is an explicit static property
- None addresses dynamic placement

Fabric: A Platform for Secure Distributed Computation and Storage

Jed Liu, Michael D. George, K. Vikram
Xin Qi, Lucas Wayne, Andrew C. Myers
{liujed, mdgeorge, kvikram, qxinqin, lrw, andru}@cs.cornell.edu
Department of Computer Science
Cornell University
4130 Upson Hall, Ithaca NY

... and language for building secure distributed systems. It is a decentralized system that allows heterogeneous domains to securely share both information and compute mutual distrust. Its high-level programming model supports data-shipping and functional computation: both computation and information flow are explicitly tracked and controlled. Fabric provides a rich, Java-like object model, labeled with confidentiality and integrity policies, and supports secure communication, time and resource accounting, and secure computation. It enforces security policies and does not offer general, principled techniques for implementing the functionality of these systems while also satisfying their security and privacy requirements. This lack motivates the creation of Fabric, a platform for building secure distributed information systems. It is particularly difficult to build secure federated systems, which integrate information and computation from independent administrative domains—each domain has policies for security and privacy, but does not fully trust other domains to enforce them. Integrating information from different domains is important because it enables new services and capabilities. To illustrate the challenges, consider the scenario of two medical institutions that want to securely and quickly share patient information. This goal is important: according to a 1999 Institute of Medicine study, at least 44,000 deaths annually result from medical errors, with incomplete patient information identified as a leading cause [25]. However, automated sharing of patient data poses difficulties. First, the security and privacy policies of the two institutions must be satisfied (as mandated by HIPAA [22] in the U.S.), restricting which information can be shared or modified by the two institu-

Adam Chlipala
MIT CSAIL
adamc@csail.mit.edu




Abstract

The World Wide Web has evolved gradually from a document delivery platform to an architecture for distributed programming. This largely unplanned evolution is apparent in the set of interconnected languages and protocols that any Web application must manage. This paper presents UrWeb, a domain-specific, statically typed functional programming language with a much simpler model for programming modern Web applications. UrWeb's model is unified, where programs in a single programming language are compiled to other “Web standards” languages as needed; supports novel kinds of encapsulation of Web-specific state; and exposes simple concurrency, where programmers can reason about distributed, multithreaded applications via a mix of transactions and cooperative preemption. We give a tutorial introduction to the main features of UrWeb and discuss the language implementation and the production Web applications that use it.

for network communication, and on a language or API like SQL for storing persistent, structured data on servers. Code fragments in these different languages are often embedded within each other in complex ways, and the popular Web development tools provide little help in catching inconsistencies.

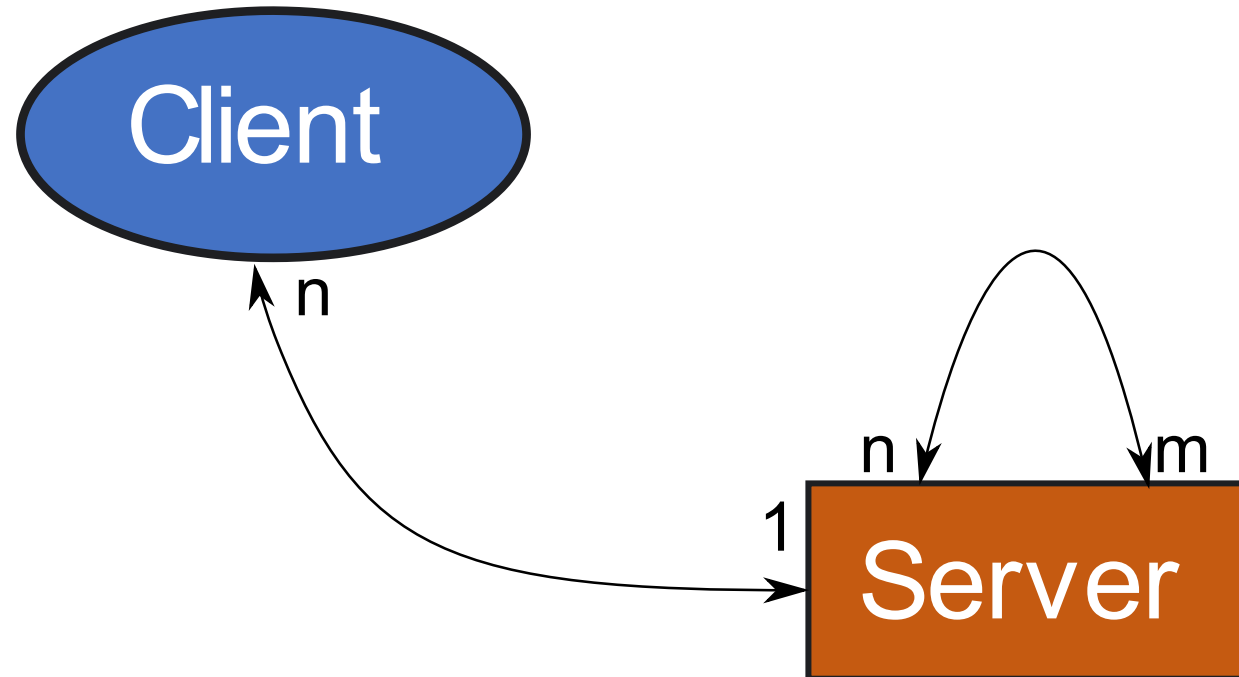
These complaints are not new, nor are language-based solutions. The Links project [7, 11] pioneered the “tierless programming” approach, combining all the pieces of dynamic Web applications within one statically typed functional language. More recent designs in the mainstream reap some similar benefits, as in Google’s Web Toolkit [3] and Closure [5] systems, for adding compilation on top of Web-standard languages; and Microsoft’s LINQ [27], for type-safe querying (to SQL databases and more) within general-purpose languages.

Such established systems provide substantial benefits to Web programmers, but there is more we could ask for. This paper focuses on a language design that advances the state of the art by ad-



Achieving Safe Dynamic Placement
with Placement Types

What are Placement Types?



What are Placement Types?

```
// The architecture is:
@peer type Client <: { type Tie <: Single[Server] }
@peer type Server <: { type Tie <: Multiple[Client] with Multiple[Server] }

// All clients have this String variable
var message: String on Client = "Greetings Earthlings"

// All servers can fetch the messages of their connected clients
def hashMessages(seed: Int): List[String] on Server =
  on[Client].run.capture(seed) {
    hash(message, seed)
  }.asLocalFromAll
```

A Principled Approach to Placement Safety

Re-Interpreting Type System Features as Placements

Interpreting Types as Placements

Type System

Placement System

Types
Int, Boolean, String

Placement Types
Client, Server, Database

Param. Polymorphism
Map[K,V], Set[T], Either[S,T]

Place-Polymorphic Modules
LeaderElection[P], DistHashMap[P]

?

Dynamic Placement
DB/Cache, CPU/GPU, LocalFS/RemoteFS

[OOPSLA 2018]



[ECOOP 2019]



Interpreting Types as Placements

Type System

Placement System

Types
Int, Boolean, String


Placement Types
Client, Server, Database

Param. Polymorphism
Map[K,V], Set[T], Either[S,T]

Place-Polymorphic Modules
LeaderElection[P], DistHashMap[P]

Union Types
Peano Numbers, Linked Lists, Options

Dynamic Placement
DB/Cache, CPU/GPU, LocalFS/RemoteFS

[OOPSLA 2018]  [Interpreting System Development with Solitaires](#)

[ECOOP 2019] [Multi-Kind Modules](#)

[OOPSLA 2023] [Type-Safe Dynamic Placement with First-Class Placed Values](#)



Union Placement Types
Statically Capture
Placement Uncertainty



Solution: Union Placement Types

- First-class placed type: **Data at (P | Q)**
- Static uncertainty and dynamic certainty
- Eliminate runtime checks when access is safe



Solution: Union Placement Types

- Remote references inhabit placement unions
- Introduced with `remote ref`

```
def readFromCache(k: Key): (Data at Cache) on Cache =  
  remote ref query(k)
```

- Eliminated with `deref` – typechecks if the architecture allows it

```
readFromCache("splash23").deref : Future[Value]
```

SAFE



Solution: Union Placement Types

- Introduced via subsumption $T \text{ at } P <: T \text{ at } (P \mid Q)$

```
def readFromCache(k: Key): Data at (DB | Cache) on Cache =  
  remote ref query(k)
```

- Eliminated with `toEither[P, Q]`

```
data.toEither[Cache, DB] : Either[Data at Cache, Data at DB]
```



```
toString()  
tsWith("file://") {  
  never happen  
oid.os.FileUriExpose
```

Evaluation: URLs as
References

Evaluation: URLs as References

```
20  ✓    public static void getBackgroundImage(ImageView imageView) {
21        String backgroundUrl = PreferenceData.BACKGROUND_IMAGE.getValue(imageView.getContext());
22
23        if (backgroundUrl != null && backgroundUrl.length() > 0) {
24            if (backgroundUrl.startsWith("http"))
25                Glide.with(imageView.getContext()).load(backgroundUrl).into(imageView);
26            else if (backgroundUrl.contains(":/")) {
27                if (backgroundUrl.startsWith("content://")) {
28                    String path = Uri.parse(backgroundUrl).getLastPathSegment();
29                    if (path != null && path.contains(":"))
30                        path = "/storage/" + path.replaceFirst(":", "/");
31                    else path = Uri.parse(backgroundUrl).getPath();
32
33                    //      "a haiku"
34                    //I don't like storage
35                    //I'm sorry, poor developer
36                    //this is all my fault
37                    //      - james fenn, 2018
38
39                    Glide.with(imageView.getContext()).load(new File(path)).into(imageView);
40                } else Glide.with(imageView.getContext()).load(Uri.parse(backgroundUrl)).into(imageView);
41            } else Glide.with(imageView.getContext()).load(new File(backgroundUrl)).into(imageView);
42        }
43    }
```

<https://github.com/fennifith/Alarmio/blob/main/app/src/main/java/me/jfenn/alarmio/Utils/ImageUtils.java>

Evaluation: URLs as References

```
20  ✓ public static void setBackgroundImage(ImageView imageView) {
21      String backgroundUrl = PreferenceData.BACKGROUND_IMAGE.getValue(imageView.getContext());
22
23      if (back
24          if
25
26          else
27              //      "a haiku"
28              //I don't like storage
29              //I'm sorry, poor developer
30              //this is all my fault
31              //
32              //      - james fenn, 2018
33
34
35
36      //
37      //      - james fenn, 2018
38
39      Glide.with(imageView.getContext()).load(new File(path)).into(imageView);
40      } else Glide.with(imageView.getContext()).load(Uri.parse(backgroundUrl)).into(imageView);
41      } else Glide.with(imageView.getContext()).load(new File(backgroundUrl)).into(imageView);
42
43
```

<https://github.com/fennifith/Alarmio/blob/main/app/src/main/java/me/jfenn/alarmio/Utils/ImageUtils.java>

Evaluation: URLs as References

```
20  ✓ public static void getBackgroundImage(ImageView imageView) {
21      String backgroundUrl = PreferenceData.BACKGROUND_IMAGE.getValue(imageView.getContext());
22
23      if (backgroundUrl != null && backgroundUrl.length() > 0) {
24          if (backgroundUrl.startsWith("http")) ←
25              Glide.with(imageView.getContext()).load(backgroundUrl).into(imageView);
26          else if (backgroundUrl.contains(":/")) { ←
27              if (backgroundUrl.startsWith("content://")) { ←
28                  String path = Uri.parse(backgroundUrl).getLastPathSegment();
29                  if (path != null && path.contains(":")) ←
30                      path = "/storage/" + path.replaceFirst(":", "/");
31                  else path = Uri.parse(backgroundUrl).getPath();
32
33                  //      "a haiku"
34                  //I don't like storage
35                  //I'm sorry, poor developer
36                  //this is all my fault
37                  //      - james fenn, 2018
38
39                  Glide.with(imageView.getContext()).load(new File(path)).into(imageView);
40              } else Glide.with(imageView.getContext()).load(Uri.parse(backgroundUrl)).into(imageView);
41          } else Glide.with(imageView.getContext()).load(new File(backgroundUrl)).into(imageView);
42      }
43  }
```

Evaluation: URLs as References



133 out of ~3K Android apps encode dynamic placement as URLs.

```
(F-Droid) 1 if (repo.getAddress().startsWith("content://") || repo.getAddress().startsWith("file://")) {  
2     // no need to show a QR Code, it is not shareable  
3     return; }
```

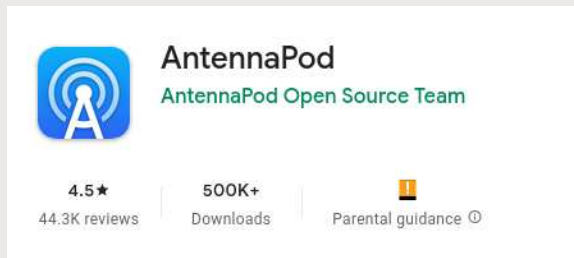

```
(VLC) 1 public Mediawrapper getMedia(Uri uri) {  
2     if ("content".equals(uri.getScheme())) return null;
```

```
(Element.io) 1 value = uri.toString()  
2 if (value.startsWith("file://")) {  
3     // it should never happen  
4     // else android.os.FileUriExposedException will be triggered.  
5     // see https://github.com/vector-im/riot-android/issues/1725  
6     return }
```

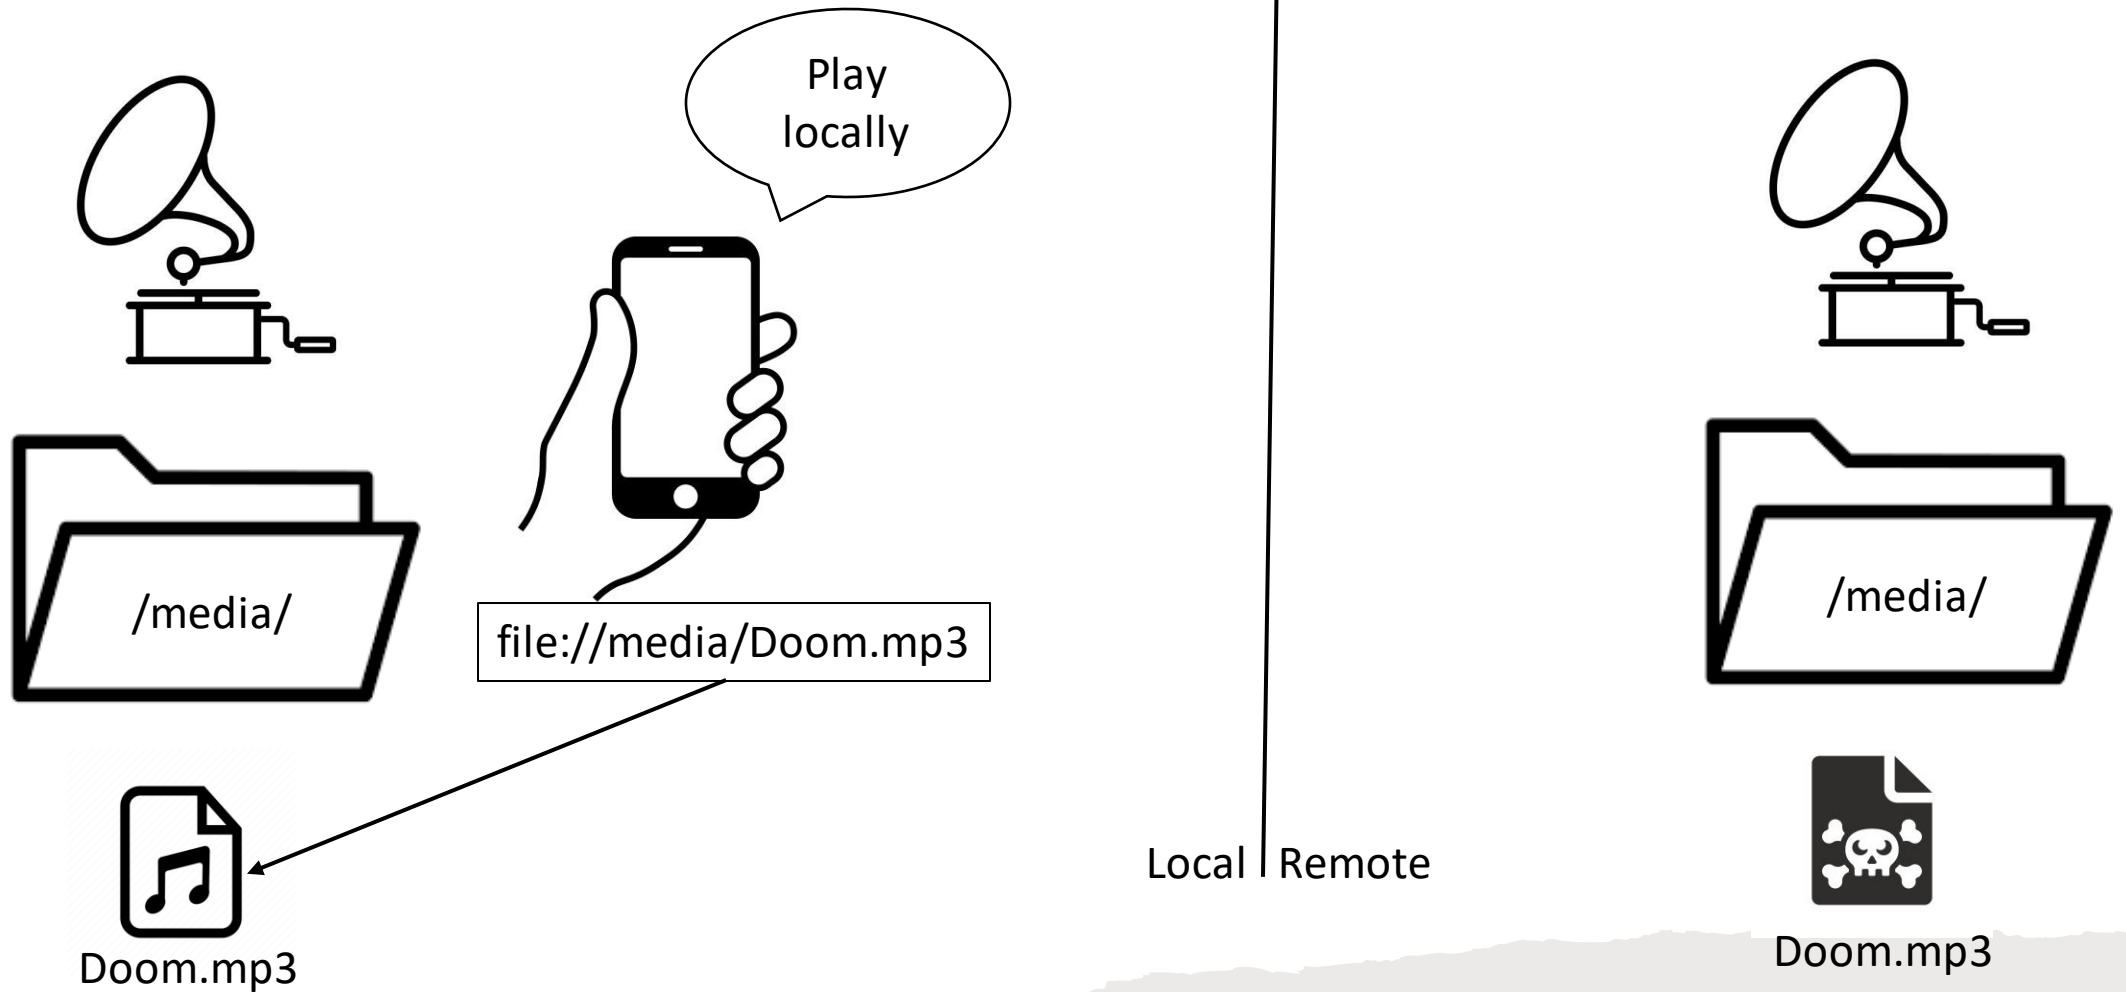


Evaluation: AntennaPod

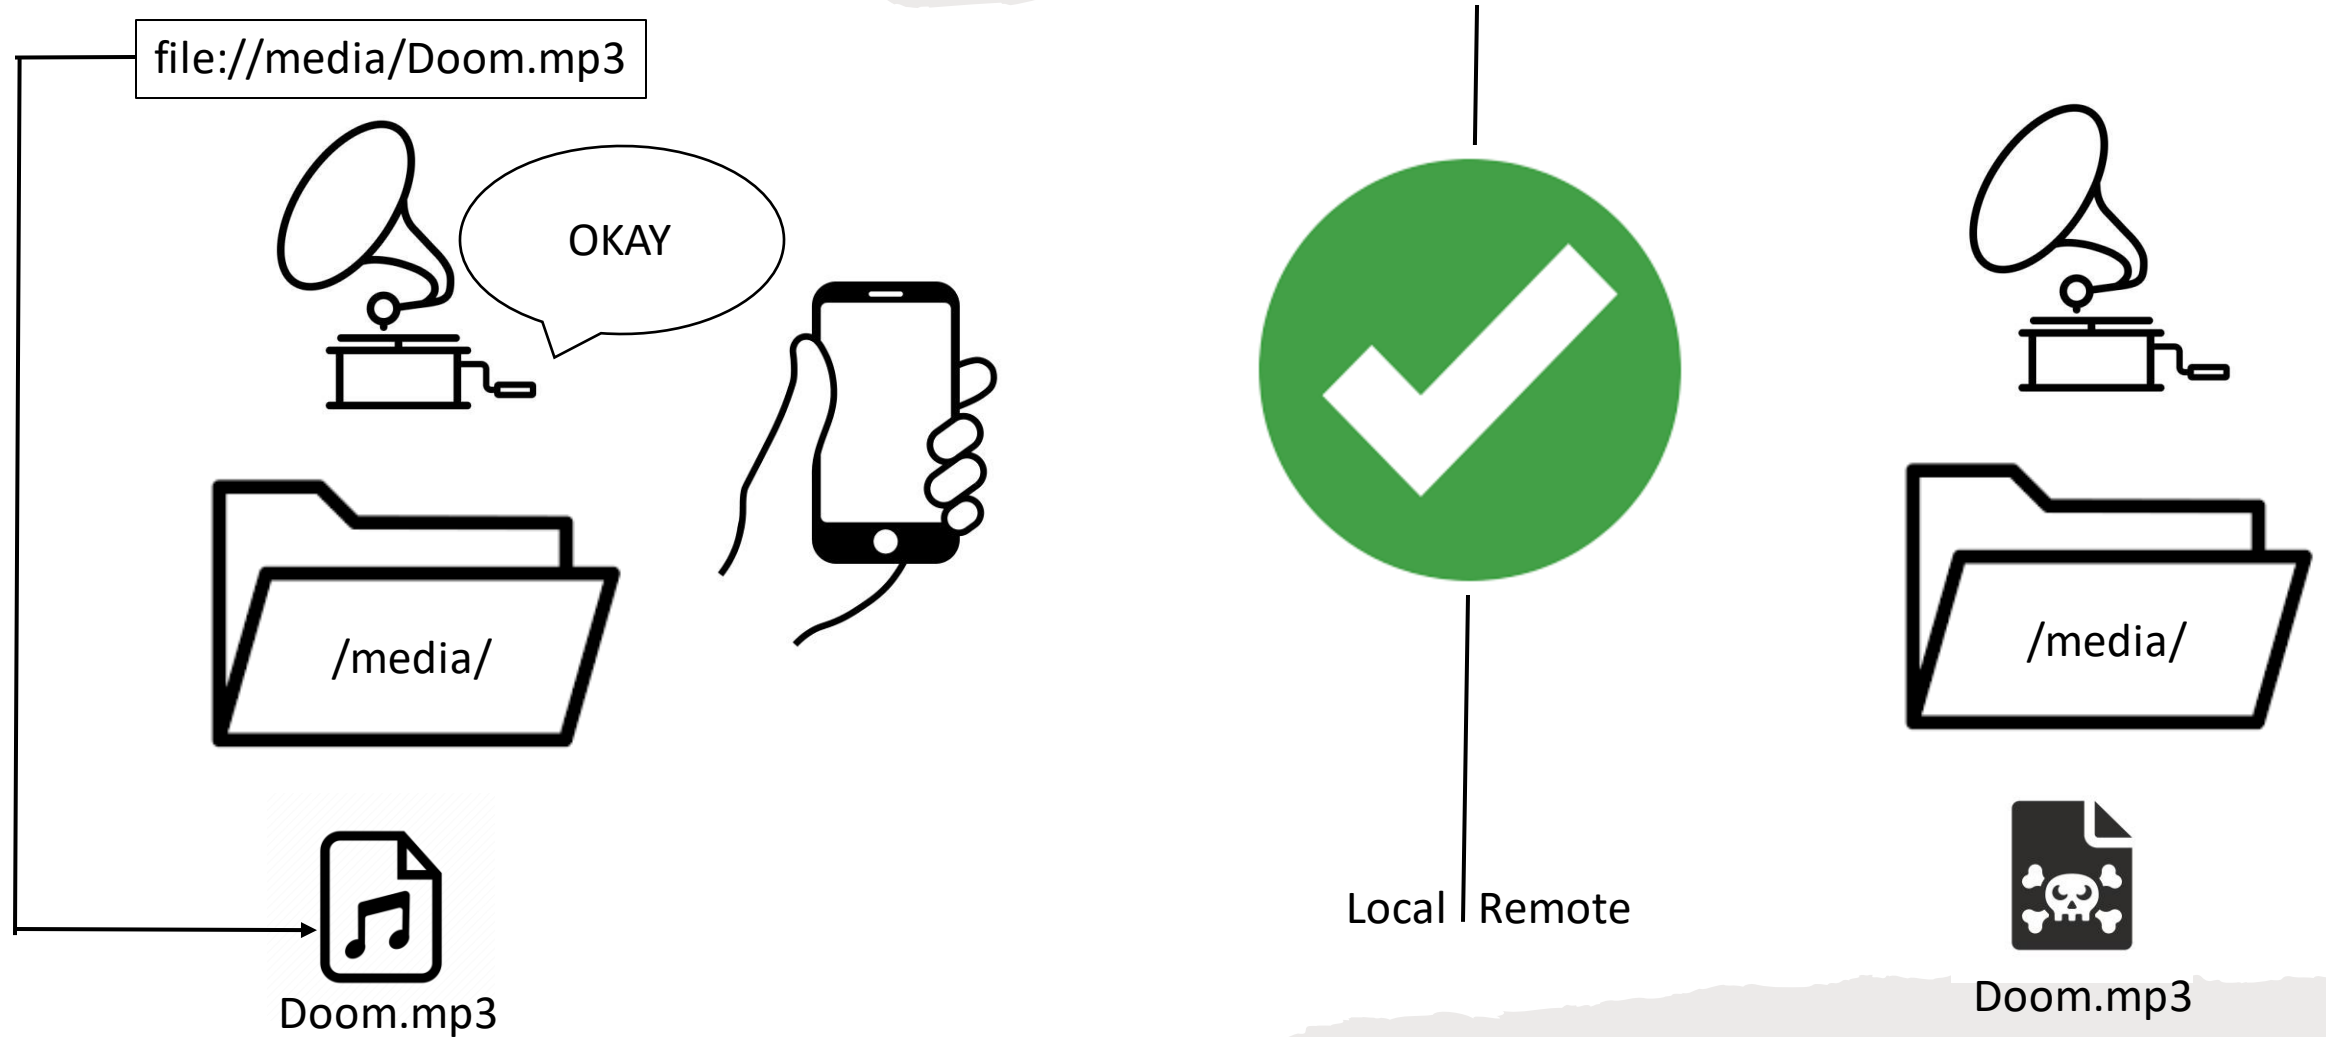
- 5.1k ★
- 1.2k forks
- 500K+ downloads on the Play Store
- ~55.5 KLoC



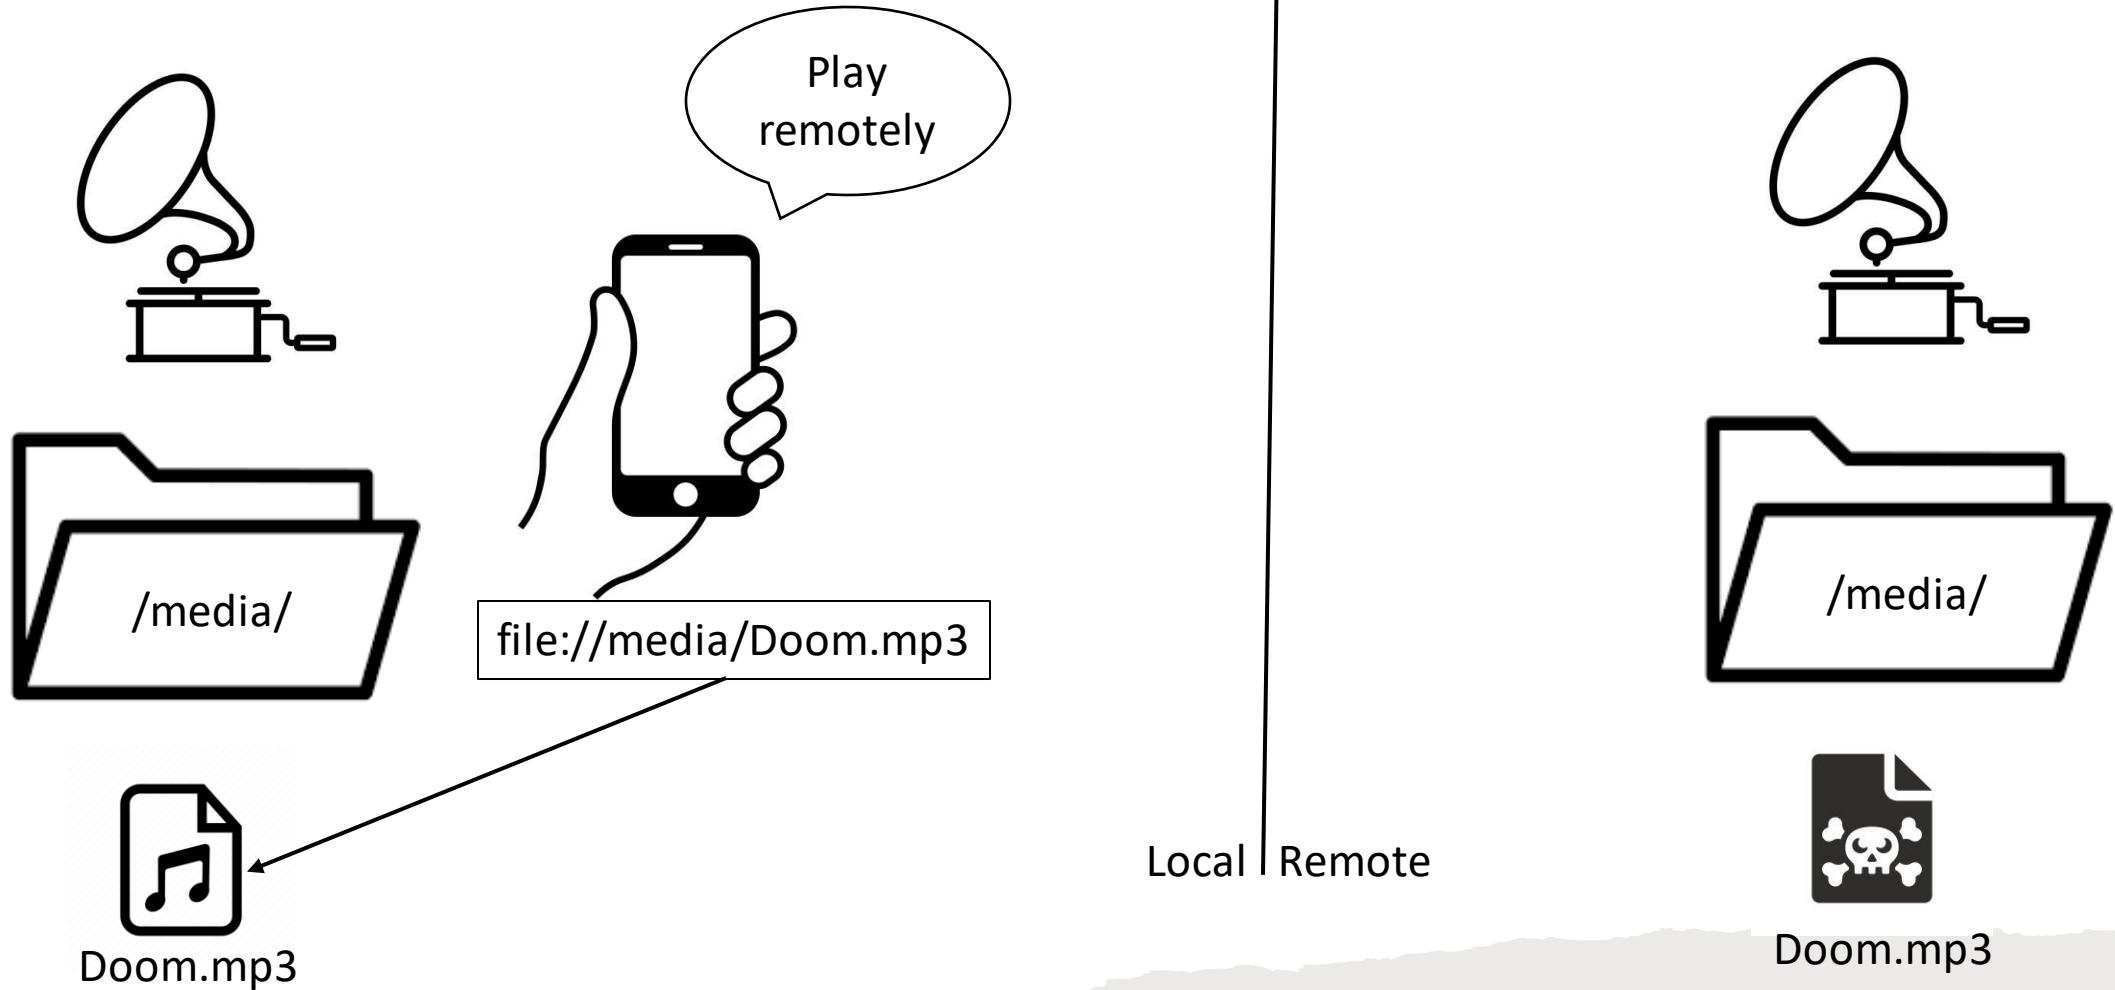
Evaluation: AntennaPod




Evaluation: AntennaPod



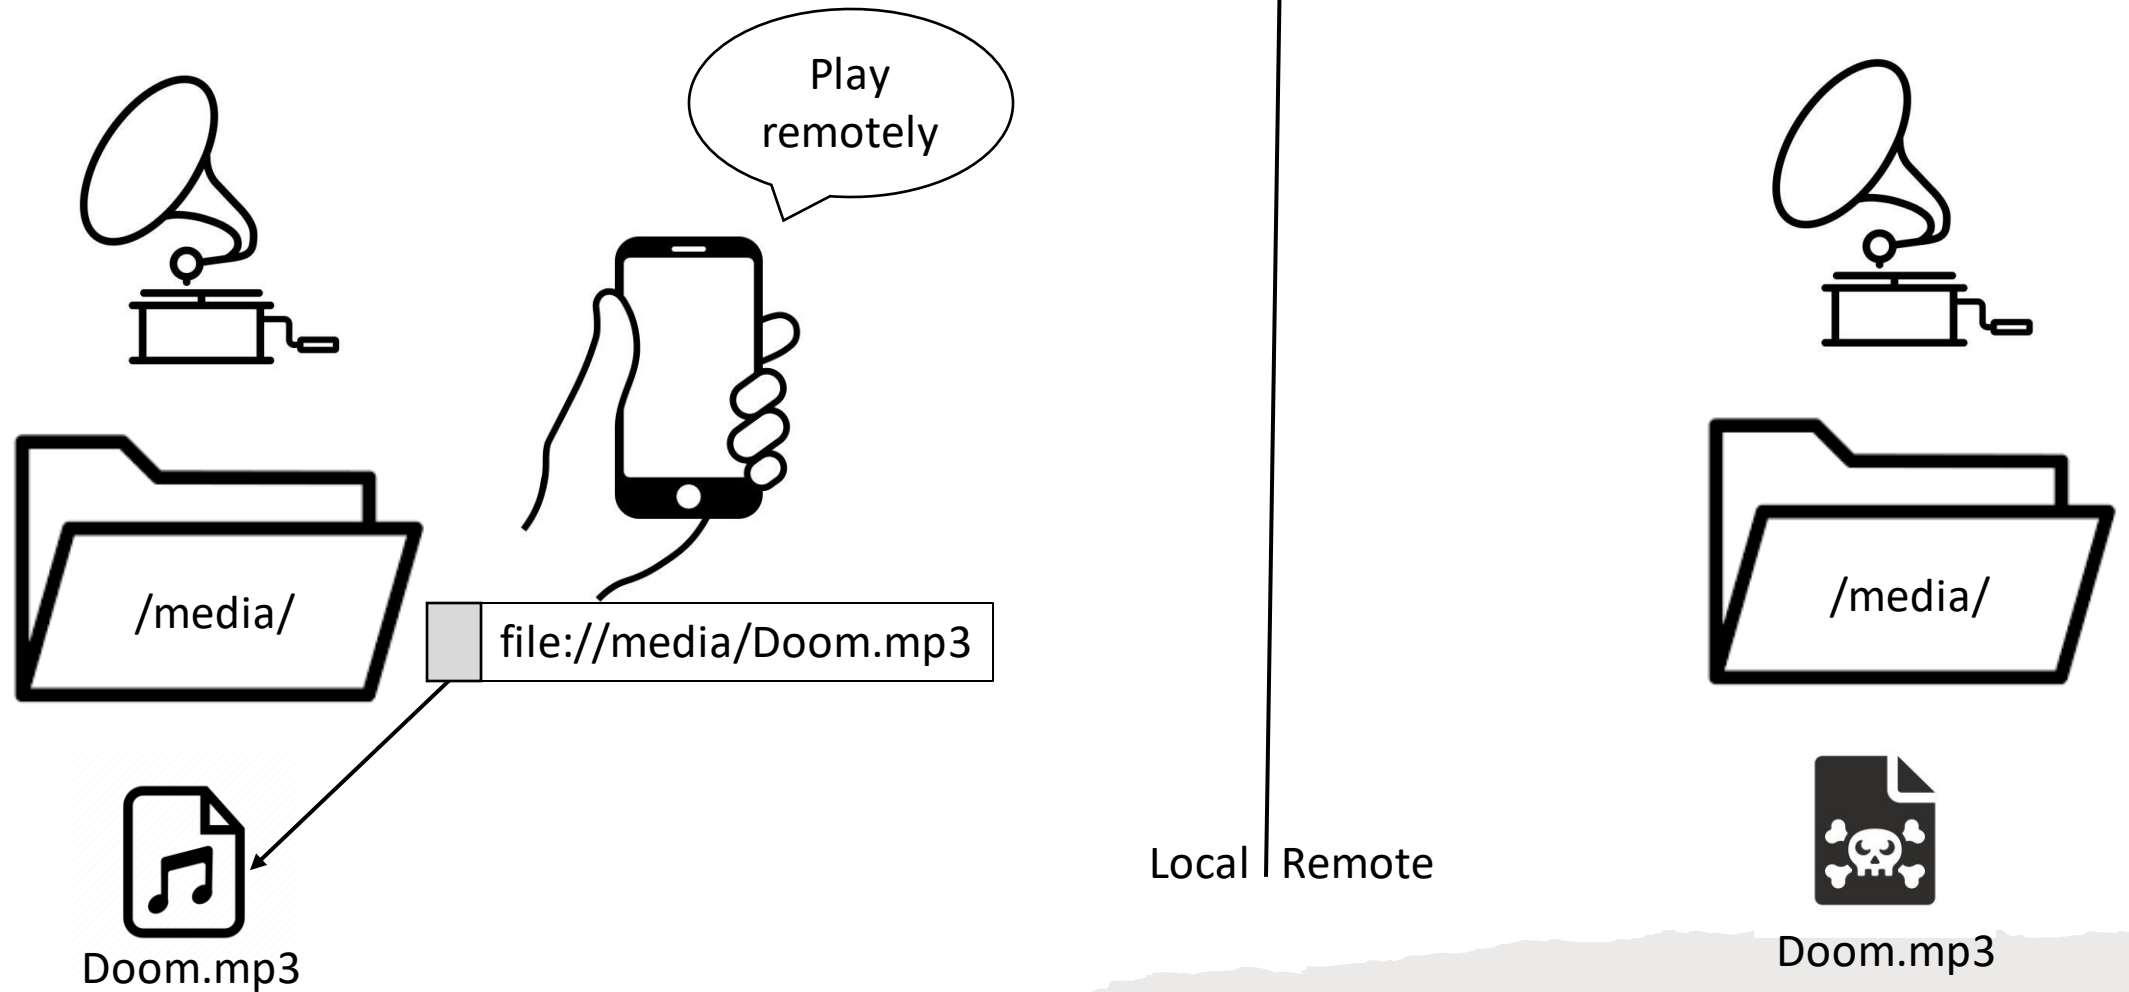
Evaluation: AntennaPod



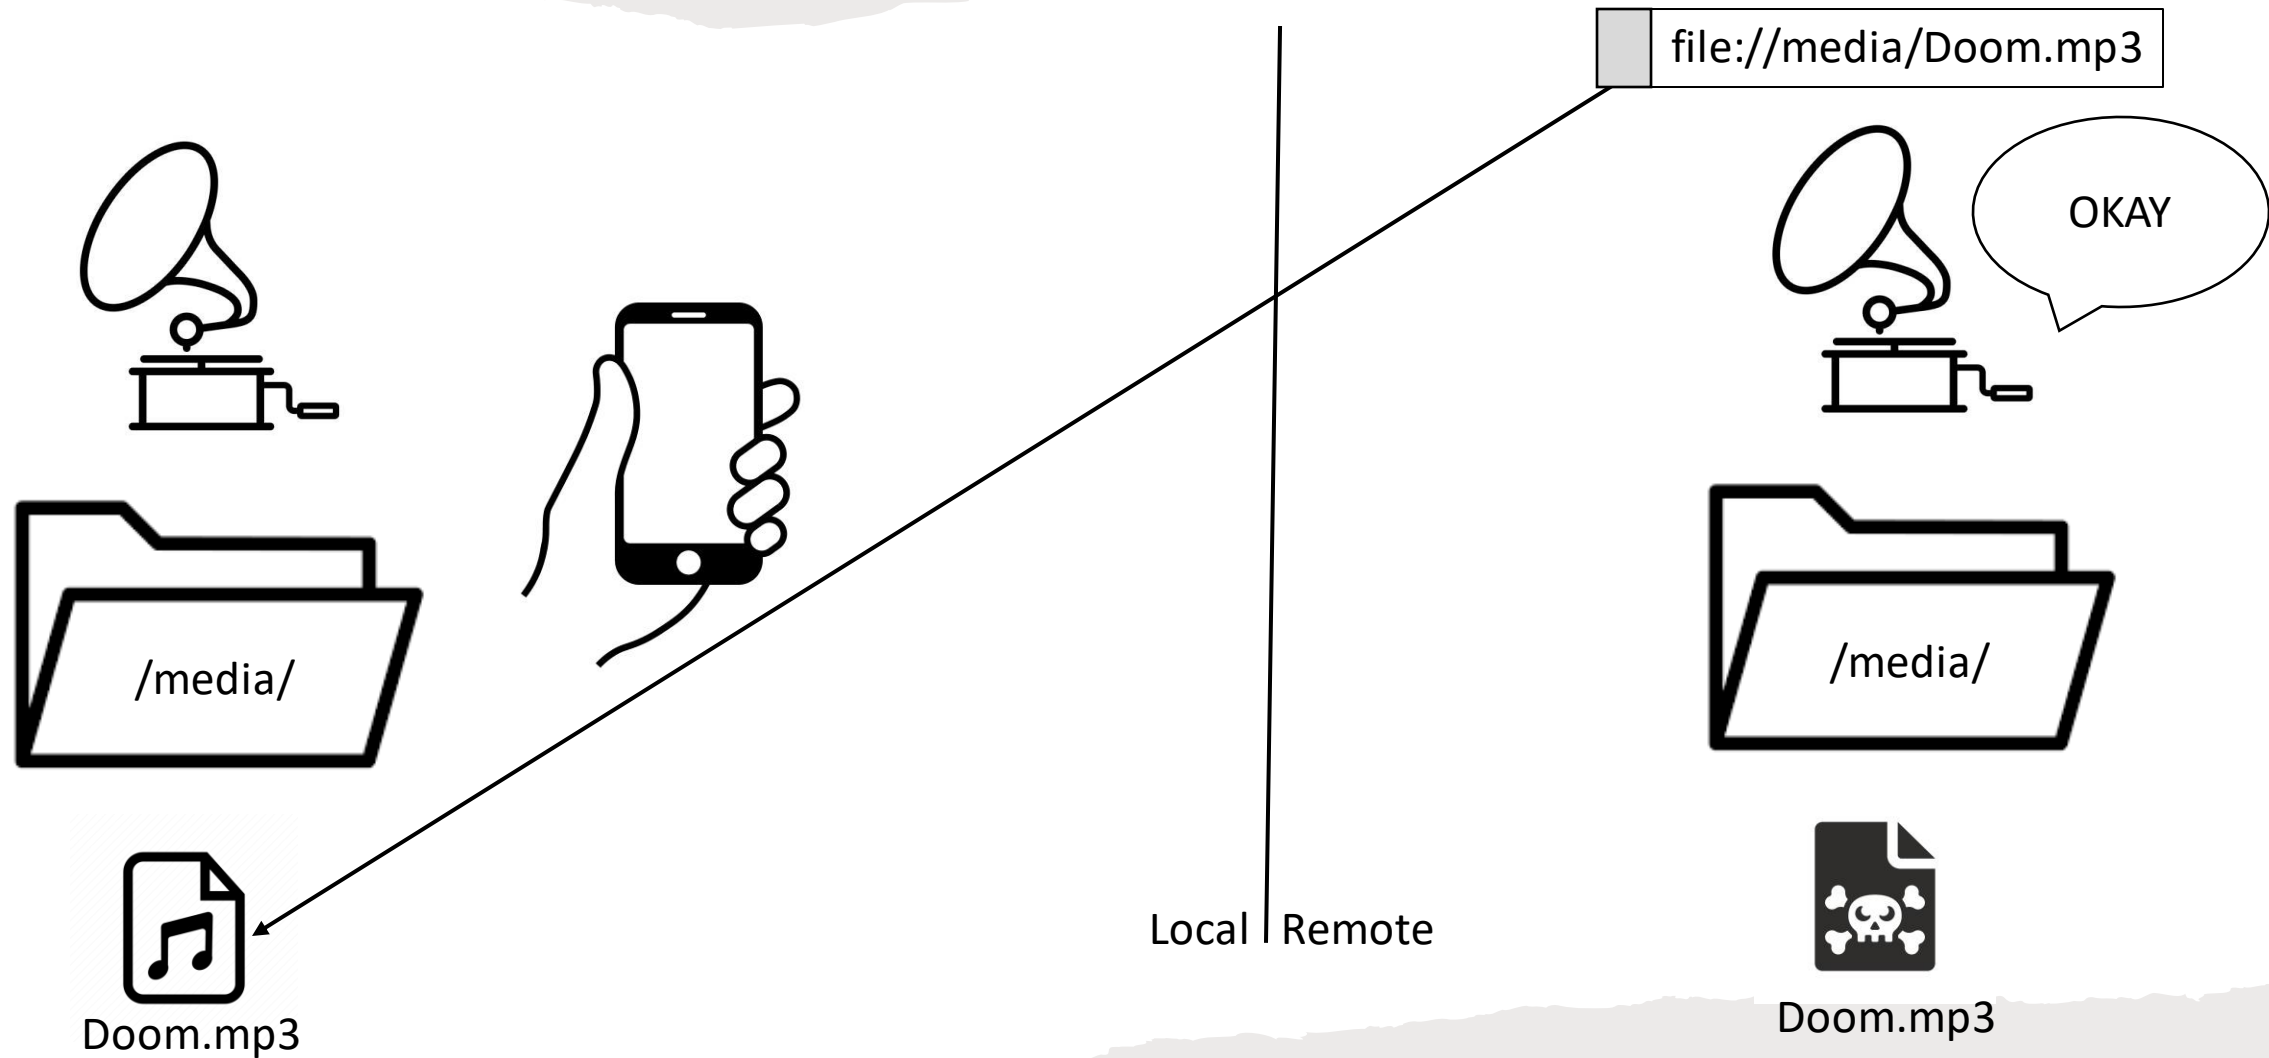
Evaluation: AntennaPod



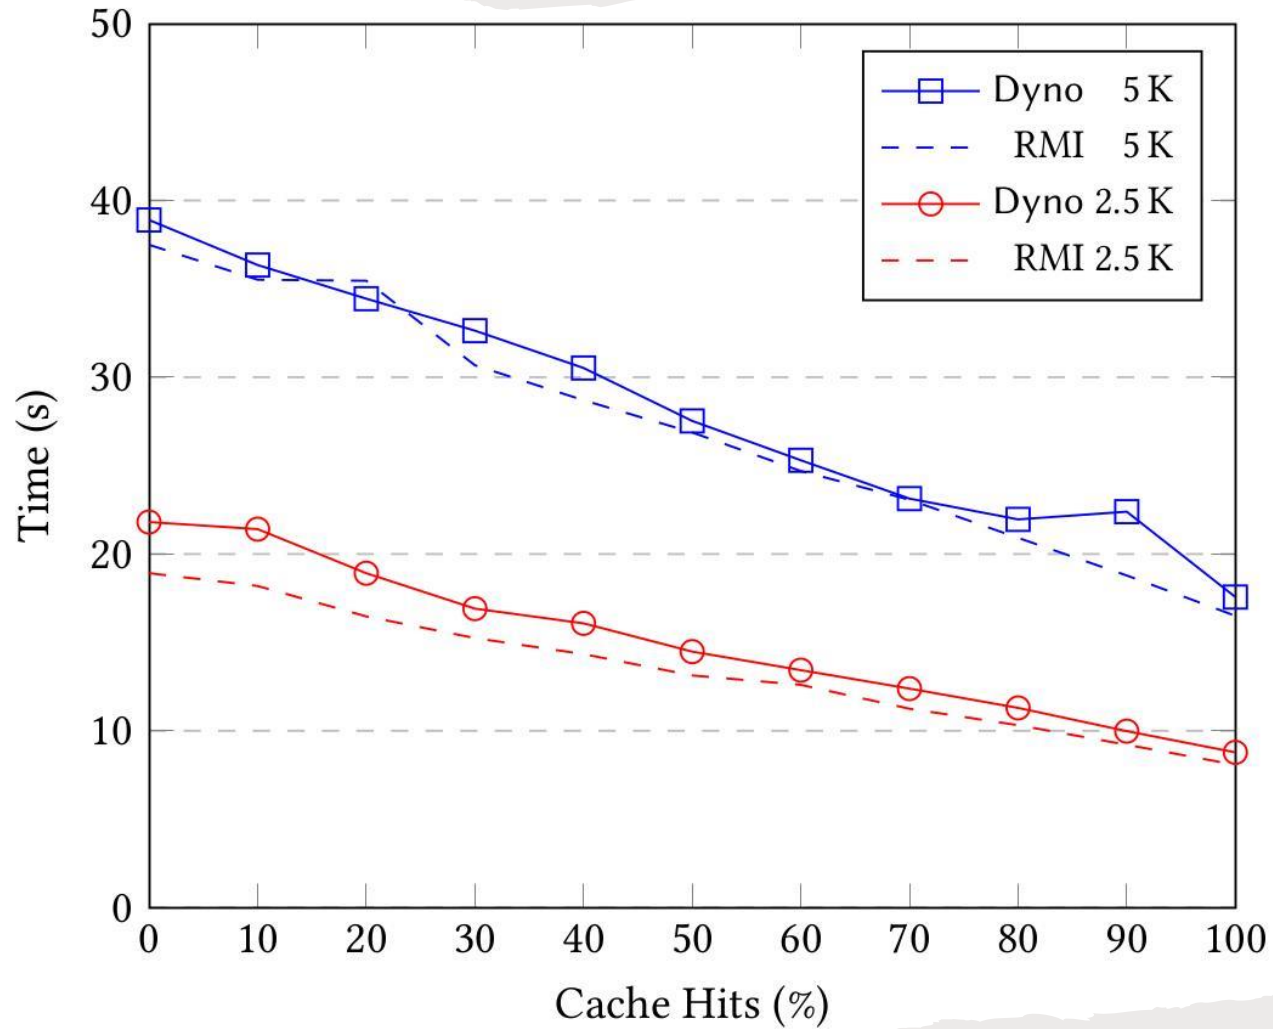
Evaluation: AntennaPod in Dyno



Evaluation: AntennaPod in Dyno



Evaluation: Performance



More goodies in the paper

- Case studies
- Comparative evaluation
- Implementation overview
- Formalization
- Theorems (and proofs in Appendix B)

Type-Safe Dynamic Placement with First-Class Placed Values

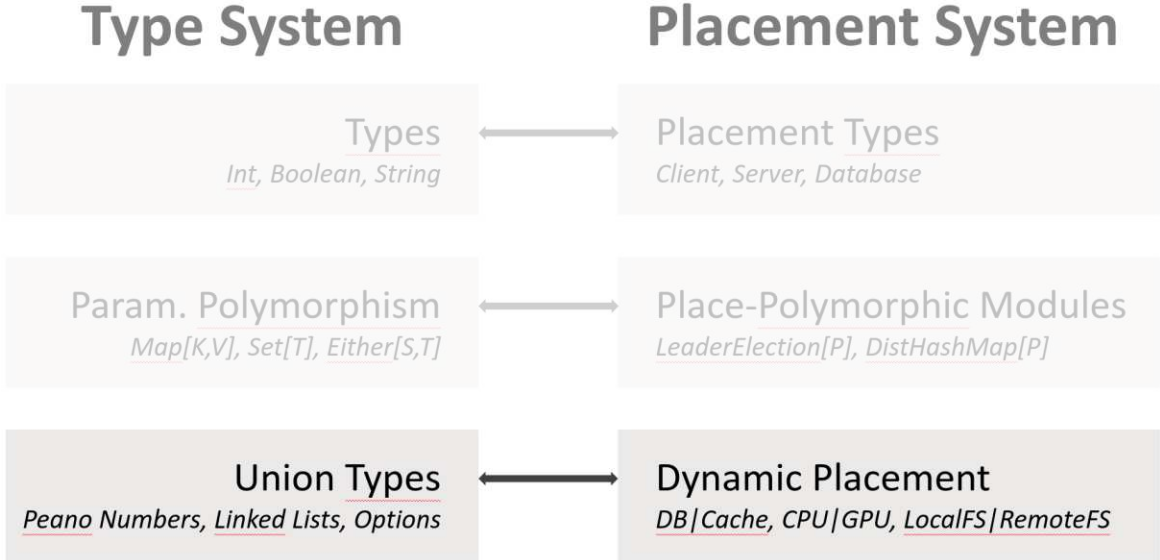

GEORGE ZAKHOUR, PASCAL WEISENBURGER, and GUIDO SALVANESCHI,
University of St. Gallen, Switzerland

Several distributed programming language solutions have been proposed to reason about the placement of data, computations, and peers interaction. Such solutions include, among the others, multitier programming, choreographic programming and various approaches based on behavioral types. These methods statically ensure safety properties thanks to a complete knowledge about placement of data and computation at compile time. In distributed systems, however, dynamic placement of computation and data is crucial to enable performance optimizations, e.g., driven by data locality or in presence of a number of other constraints such as security and compliance regarding data storage location. Unfortunately, in existing programming languages, dynamic placement conflicts with static reasoning about distributed programs: the flexibility required by dynamic placement hinders statically tracking the location of data and computation.

In this paper we present Dyno, a programming language that enables static reasoning about dynamic placement. Dyno features a type system where values are explicitly placed, but in contrast to existing approaches, placed values are also first class, ensuring that they can be passed around and referred to from other locations. Building on top of this mechanism, we provide a novel interpretation of dynamic placement as unions of placement types. We formalize type soundness, placement correctness (as part of type soundness) and architecture conformance. In case studies and benchmarks, our evaluation shows that Dyno enables static reasoning about programs even in presence of dynamic placement, ensuring type safety and placement correctness of programs at negligible performance cost. We reimplement an Android app with ~ 7 K LOC in Dyno, find a bug in the existing implementation, and show that the app's approach is representative of a common way to implement dynamic placement found in over 100 apps in a large open-source app store.

CCS Concepts: • **Software and its engineering** → **Distributed programming languages**; Domain specific languages; • **Theory of computation** → *Distributed computing models*.

Conclusion



Data at (P | Q)