

Blueprint: A Toolchain for Highly-Reconfigurable Microservices

Vaastav Anand, Deepak Garg, Antoine Kaufmann, Jonathan Mace
Max Planck Institute for Software Systems (MPI-SWS)

Abstract

Researchers and practitioners care deeply about the performance and correctness of microservice applications. To investigate problematic application behavior and prototype potential improvements, researchers and practitioners *experiment* with different designs, implementations, and deployment configurations. We argue that a key requirement for microservice experimentation is the ability to rapidly reconfigure applications and to iteratively Configure, Build, and Deploy (CBD) new variants of an application that alter or improve its design. We focus on three core experimentation use-cases: (1) updating the design to use different components, libraries, and mechanisms; (2) identifying and reproducing problematic behaviors caused by different designs; and (3) prototyping and evaluating potential solutions to such behaviors. We present Blueprint, a microservice development toolchain that enables rapid CBD. With a few lines of code, users can easily reconfigure an application’s design; Blueprint then generates a fully-functioning variant of the application under the new design. Blueprint is open-source and extensible; it supports a wide variety of reconfigurable design dimensions. We have ported all major microservice benchmarks to it. Our evaluation demonstrates how Blueprint simplifies experimentation use-cases with orders-of-magnitude less code change.

ACM Reference Format:

Vaastav Anand, Deepak Garg, Antoine Kaufmann, Jonathan Mace. 2023. Blueprint: A Toolchain for Highly-Reconfigurable Microservices. In *ACM SIGOPS 29th Symposium on Operating Systems Principles (SOSP ’23)*, October 23–26, 2023, Koblenz, Germany. ACM, New York, NY, USA, 30 pages. <https://doi.org/10.1145/3600006.3613138>

1 Introduction

Modern cloud applications are increasingly developed as suites of loosely-coupled microservices [17]. The microservice architectural approach decomposes applications into

smaller, modular services that communicate over the network. As a result of their success in enabling large-scale and continually-evolving applications, microservices have become ubiquitous among internet companies including Twitter [30], Netflix [16], Uber [29], and others [17].

Microservices are large and complex applications, composed of multiple pieces including frameworks, backends, and libraries. For any application, there are many possible designs, each with its own set of performance properties and issues. As a result, researchers are highly interested in studying, measuring, and improving the performance and correctness guarantees of microservice systems, and developing solutions to potential unwanted behavior when they arise. A salient example of unwanted behaviour are *emergent phenomena* [48] of microservice systems which include cascading failures [52, 70, 75], tail latency effects [18], cross-system consistency [3], and metastable failures [15, 33], among others. By analyzing these behaviors, researchers hope to improve the performance and correctness guarantees of microservice systems, and develop solutions to potential unwanted phenomena [15, 33, 43, 55, 61].

Towards this goal, researchers, both in academia and industry, need to be able to perform three basic actions: (i) reconfigure applications with new features, backends, and libraries to improve their performance and add new features; (ii) reproduce and discover potentially problematic emergent phenomena of applications; and (iii) develop and evaluate solutions on canonical microservice systems. The central requirement of these use-cases is for researchers to *easily explore the design space of microservices*, allowing them to move between different implementations and deployment configurations of the same application quickly and easily.

Enabling design space exploration is difficult for several reasons. First, the design space of microservice systems is enormous. Application design, configuration, and deployment choices vary along several dimensions: specific patterns in the flow of application logic (*e.g.* the presence of concurrent or asynchronous execution branches); the inclusion of particular features, components, or middleware (*e.g.* replicated services, autoscaling, and sharding); and component instantiations and their corresponding configuration (*e.g.* the specific RPC framework used and its timeout and retry mechanisms). Given the lack of standardization across these axes, it is not immediately clear how to support all possible application design dimensions and all possible instantiations for a given dimension in a systematic manner.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

SOSP ’23, October 23–26, 2023, Koblenz, Germany

© 2023 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0229-7/23/10.

<https://doi.org/10.1145/3600006.3613138>

Second, existing microservice systems, both experimental and production-grade, are implemented as point solutions in this vast design space and are not designed to be reconfigurable or extensible. Thus, the entire burden for moving from one implementation to another falls on the researcher as they have to make deep modifications to microservice applications to fulfill their use-case. A high-level design change, such as replacing the RPC framework, typically translates into thousands of LoC of source-level implementation changes, dispersed across the many processes, components, and backends that comprise the microservice application (cf. §3.1). Implementing a design change is time-consuming, error-prone, and difficult, as it requires understanding the application at the source-code level; yet microservice implementations tightly couple concerns at the source-code level, *i.e.* application logic directly binding to libraries and middleware. Thus, changes that are conceptually simple – *e.g.* replacing an existing RPC framework with a different one – can require wide-ranging and complex manual modifications to many components. Most prior research work on microservices had no alternative to expending this developer effort, and consequently the majority of works deploy and evaluate solutions for only a single application (78%, cf. §B); anecdotally, the developer burden is the limiting factor.

Due to the large design space of microservice applications, there is a significant concern about the generalizability of research results derived from only a single application. A solution may have implicit dependencies on particular application design choices, or worse, it may be an application-specific solution that does not apply broadly at all. For example, Sage [25] assumes synchronous RPCs and Tprof’s [34] layer4 grouping assumes non-combinatorial explosion when grouping requests by visited services’ execution order; both of which do not hold at Meta [35]. Under normal circumstances, a lack of broad evaluation would be considered a benchmarking crime [31]. However, for microservices it is the prevailing norm.

The goal of this work is to enable researchers to easily reconfigure microservice applications. Our solution, Blueprint, is a microservice development toolchain designed for rapidly *Configuring, Building, and Deploying* (CBD) microservices. Blueprint enables its users to easily *mutate* the design of an application and generate a fully-functional variant that incorporates their desired changes.

The key insight of Blueprint is that the design of a microservice application can be decoupled into (i) the application-level *workflow*, (ii) the underlying *scaffolding* components such as replication and auto-scaling frameworks, communication libraries, and storage backends, and (iii) the concrete *instantiations* of those components and their configuration. An application written using Blueprint avoids tightly coupling these concerns. Instead, these *design aspects are concisely declared* by the user using Blueprint’s abstractions,

namely, the *workflow spec* and the *wiring spec*. Blueprint’s compiler combines these two abstractions to *automatically generate the system*. Changing any given aspect only requires the developer to revisit its declaration in Blueprint’s abstractions and not the generated implementation. Moreover, Blueprint eliminates duplicated effort – scaffolding and instantiation logic are implemented once and integrated as Blueprint compiler extensions, to enable Blueprint to automatically change existing Blueprint applications with minimal effort. Concretely, using Blueprint, users can:

- mutate off-the-shelf microservice applications (*e.g.* an open-source benchmark) with just a few LoC, to swap an instantiation (*e.g.* RPC framework), enable or disable scaffolding (*e.g.* replication), or change backends (*e.g.* database).
- develop new application workflows and generate runnable systems. Instead of binding workflow code to specific scaffolding and instantiations (*e.g.* the choice of RPC framework), those are declared separately with 10s of LoC, and incorporated by Blueprint at compile time.
- introduce support for new instantiations (*e.g.* an experimental RPC framework) or scaffolding concepts (*e.g.* geo-replication) and transparently apply them to existing applications; these are implemented as compiler extensions that are independent and agnostic to specific application workflows.

We have implemented Blueprint in 11k LoC of Go and ported all major available microservice benchmarks to Blueprint (15k LoC), including the DeathStar benchmark [26], TrainTicket [76], and SockShop [47]. Our evaluation demonstrates the ease with which Blueprint enables changes to the design and features of an application; we reproduce known interesting behaviour from a number of prior works; and show that Blueprint is easily extensible with new features that can be reused across applications.

2 Microservice Design Space

The space of possible designs for microservice applications is enormous, and while there may be guiding high-level design principles, every application differs substantially from the next. In this section, we briefly elaborate three important dimensions for the design of a microservice application. These axes are useful both for motivating Blueprint’s use cases (§3) and as insights into Blueprint’s design (§4).

A microservice application’s design can be principally decomposed along three major dimensions: the **application-level workflow**; the lower-level **scaffolding** infrastructure on which the workflow executes; and **instantiations** for different infrastructure implementations.

Application-Level Workflow. Microservice applications vary widely in terms of their application-level logic and end-to-end flow of executions through the system’s different components. Recent open-source microservice benchmarks cover

diverse domains, such as e-commerce apps [28, 47, 71, 72], social networks [26], reservation systems [26, 76] and many others [2, 6, 26, 36, 66]. These applications differ in the number of services and APIs they use internally from only a few (SockShop [47]) to dozens (TrainTicket [76]). Even applications with similar high-level goals (e.g. TrainTicket [76] and DSB Hotel Reservation [26]) have vastly different workflows.

Scaffolding. Scaffolding refers to runtime components which are necessary for executing an application but orthogonal (and opaque) to the application’s workflow. Scaffolding includes middleware, framework, libraries and backends that provide features such as RPCs, replication, load balancing, circuit breakers, and more [11]. Scaffolding can be changed without affecting the application’s workflow and functional behavior. For example, the original DSB Social Network [26] uses one runtime instance of each service by default; however, a changed version of the application could modify its scaffolding to replicate service instances, without changing the end-to-end application workflow [42].

Instantiations. For every piece of scaffolding there may be different implementations to choose from, configuration dimensions of those implementations, and choices for configuration parameters, e.g. to enable RPC an application may use gRPC, Thrift, or some research prototype framework [39].

Overall, the choice of an application’s workflow, scaffolding, and instantiations have different implications for the application’s performance, correctness, and reliability.

3 Blueprint Use-Cases

This section outlines challenges faced by researchers today with respect to three core use-cases. Across the three use cases, we motivate a common need to iteratively Configure, Build, and Deploy (CBD) variants of microservice applications that have subtly different designs. Blueprint, which we will introduce in §4, is designed to address this need.

3.1 UC1: Mutating Applications

To investigate and experiment with changing workloads and deployment conditions, researchers may wish to *mutate* an application – i.e. change, reconfigure, or expand some aspect of its design. A mutation modifies the application along one or more of the aforementioned axes (§2). For example, changing the RPC framework from Thrift to gRPC is a mutation that only modifies instantiations; introducing replicated services and a load balancer is a mutation that modifies both scaffolding and instantiations; and refactoring the application workflow is a mutation that typically leaves scaffolding and instantiations untouched [19]. Ultimately, these changes modify the application to move its design from one point to another in the design space, creating a new variant of the application.

Ideally, users should be able to mutate an application with minimal effort. Yet today, a conceptually simple mutation

may require far-reaching and time-consuming modifications to source code and configuration. For example, switching to a different RPC framework instantiation in an application with 30 services requires modifying all of those services. Replacing one instantiation with another is difficult because interfaces offered by instantiations of the same piece of scaffolding can differ widely, and existing systems tightly couple the application’s workflow, scaffolding, and instantiations. Similarly introducing new scaffolding, such as enabling distributed tracing, entails exhaustively updating *all* services to support it; and changing an application’s workflow requires binding the new or changed code to scaffolding and instantiation interfaces.

To understand the scope of mutations to existing microservice applications, we surveyed 464 forks of popular microservices benchmarks. We found a total of 146 application variants that apply mutations that include adding tracing, removing tracing, adding replication, adding georeplication, switching RPC frameworks, and more (cf. §B.3). As an example of the cost of manual mutations, an instantiation change in DSB Social Network [26] to support Sifter [40] required 1,289 lines of manual code change and took 2 weeks to complete based on commit timestamps (cf. §6.3).

3.2 UC2: Reproducing Emergent Phenomena

Emergent phenomena, or emergent behaviors [48], are aspects of the system’s runtime behavior that are not localized to any one service or component, instead arising as the cumulative effect of interactions *between* components under a given workload and application design. Emergent phenomena encompass end-to-end performance, correctness, and reliability concerns of an application – notable examples include cascading failures [52, 70, 75], tail latency effects [18], cross-system consistency [3], laser of death [51], and metastable failures [15, 33]. See §B.1 for a detailed list of known emergent phenomena.

Left unchecked, emergent phenomena can lead to degraded service and even outages [46, 50, 51]; thus they are the focus of a range of recent work from both industry [61] and academia [15, 33, 43, 55]. In general, researchers need the ability to elicit emergent phenomena in microservice applications, to determine the conditions under which they emerge, and their effects on application behavior.

Few existing microservice systems exhibit emergent phenomena out-of-the-box, as they arise due to specific workflow, scaffolding, or instantiation choices, combined with workloads and deployment conditions. To study emergent phenomena, researchers must therefore mutate existing applications to find variants that exhibit the phenomena. For example, no out-of-the-box microservice benchmark exhibits cross-system inconsistency [43, 61] or metastable failures [15]; researchers studying these phenomena manually mutate existing microservice applications to elicit them [22, 42].

It is not straightforward to identify a design that exhibits an emergent phenomenon. Changes to the scaffolding or instantiations can affect performance and error-propagation characteristics non-linearly, making it difficult to predict the effects of even minor alterations. Moreover, certain emergent phenomena only manifest under specific workloads which are not known to developers a priori. Overall, this makes empirical exploration of the design space essential: researchers may need to mutate an application multiple times before finding a variant that clearly exhibits a given phenomenon.

3.3 UC3: Prototyping and Evaluating Improvements

Researchers often need to prototype and evaluate *improvements* to microservice applications. An improvement can include applying workflow design patterns [19], enabling novel scaffolding [42], and implementing instantiations with better properties than the existing ones [39]. Improvements typically target some performance, correctness, or reliability concern, e.g. the application’s scalability, latency, or some emergent phenomenon. For example, FIRM [55] is an improvement for mitigating SLO violations; it leverages distributed tracing scaffolding and introduces a novel orchestration scaffolding and instantiation.

To develop and evaluate improvements, it necessary to mutate existing applications to incorporate the improvement. This entails the same degree of manual overhead described for UC1 and UC2, but is exacerbated in two ways: first, development may require multiple iterations of design to converge on appropriate interfaces, potentially entailing repeated mutations over time integrating successive versions of the improvement; second, best practices for rigorous evaluation demand that any improvement should be evaluated across multiple, diverse applications [31], thus requiring effort to mutate multiple applications, not just one.

Due to the high cost of mutating applications, most research works today evaluate on only a single microservice application (78%, cf. §B.2). Consequently it is difficult to distinguish whether a proposed improvement would be similarly effective for other applications, or just for the specific application selected. Moreover, an improvement might make assumptions about an application’s design that restricts its broader applicability. For example, FIRM [55] assumes a deterministic critical path for each API, so might not apply to workflows with concurrent or branching RPC calls.

4 Design

Blueprint is a toolchain that offers first-class CBD support for microservice applications. Instead of directly implementing runnable application artifacts (e.g. code, container images, etc.), these are generated by Blueprint’s compiler. Developers are still responsible for implementing application workflows (or re-using open-source workflows), and these must adhere to Blueprint’s abstractions. Likewise developers must separately specify which scaffolding and instantiations to apply

to the workflow. Blueprint’s compiler will automatically generate the necessary artifacts (e.g. glue code, configuration, wrappers, scripts, and more) to produce a runnable variant.

Separation of Concerns. Blueprint’s key insight is that an application’s workflow, scaffolding, and instantiations are conceptually orthogonal and thus should be separable when specifying the application. The workflow of an application is independent of the specific choice of, e.g. RPC library, or the presence of particular scaffolding, e.g. replication. In today’s applications these are tightly coupled, with application code that intertwines workflow, scaffolding, and instantiations, yet the interfaces between them are narrow because they are conceptually separate concerns and little information is needed of one to specify the other. Blueprint thereby only combines workflow, scaffolding, and instantiations at compile time, thus avoiding tight-coupling or hard-coding.

Compile-time integration. Despite a clean conceptual separation between workflow, scaffolding, and instantiations, in practice these manifest in diverse ways and at different granularities, e.g. application-level libraries, sidecar processes, container images, and orchestration framework configuration. Compile-time integration thus becomes necessary for Blueprint to abstract across diverse granularities. Blueprint’s compiler encapsulates a wide range of concerns ranging from code, process, and container image generation, to templating, dynamic addressing, and configuration.

Examples. Blueprint enables users to:

- obtain variant applications by simply recompiling with different scaffolding and instantiation choices.
- mutate an application by changing as little as 1 LoC.
- change instantiations (e.g. RPC, database implementations) with as little as 1 LoC.
- develop or change workflows with less cognitive overhead, since workflow logic is not coupled with scaffolding or instantiations.
- integrate new instantiations and scaffolding concepts as one-shot compiler plugins, reusable by any existing or future application. Implementing a plugin does not require knowledge of or compatibility with other plugins.

4.1 Blueprint Applications

A Blueprint application consists of two key abstractions. The *workflow spec* abstraction is the implementation of the application’s workflow. The *wiring spec* abstraction declares the scaffolding and instantiations to apply to the workflow. We describe each in detail.

Workflow Spec. The basic building block of a workflow is Blueprint’s *service* abstraction: developers can declare an interface for the service with typed methods, and provide an implementation of those methods. Blueprint currently supports Go. Fig. 1 defines a `ComposePost` service from the DSB Social Media application [26] that enables callers to

```

1 type ComposePostService interface {
2   ComposePost(ctx context, userID int64, text postContent) error
3 }
4 type ComposePostImpl struct {
5   postStorageService PostStorageService
6   userService UserService
7 }
8 func NewComposePostImpl(ps PostStorageService, us UserService) *
9   ComposePostService {
10  return &ComposePostImpl{ps, us}
11 }
12 func (c *ComposePostImpl) ComposePost(ctx context, userID int64, text
13   postContent) error {
14   creator, err := c.userService.GetUser(ctx, userID)
15   post := Post{Creator: creator, Text: text}
16   return c.postStorageService.StorePost(ctx, post)
17 }

```

Fig. 1. Workflow Spec for DSB Social Network ComposePost.

```

1 type Cache interface {
2   Put(key []byte, value []byte) error
3   Get(key []byte) ([]byte, error)
4 }

```

Fig. 2. Built-in interface for a *cache* backend component.

```

1 normal_deployer : Modifier = Docker()
2 rpc_server : Modifier = GRPCServer()
3 tracer : Tracer = ZipkinTracer()
4 tracerModifier: Modifier = TracerModifier(tracer=tracer)
5 server_modifiers : List[Modifier] = [rpc_server, normal_deployer,
6   tracerModifier]
7 post_cache := Memcached()
8 post_db = MongoDB()
9 user_db = MongoDB()
10 us = UserServiceImpl(user_db).WithServer(server_modifiers)
11 ps = PostStorageServiceImpl(post_cache, post_db).WithServer(
12   server_modifiers)
13 c1 = Container(ps, post_cache)
14 cs = ComposePostServiceImpl(ps, us).WithServer(server_modifiers)

```

Fig. 3. Wiring Spec for three dependent DSB services. Zipkin tracing is enabled for all services; they are deployed in Docker containers; and communicate using gRPC. Cache and database instantiations are Memcached and MongoDB respectively.

upload new social media posts. Method implementations can be arbitrary and import arbitrary libraries.

Blueprint’s *backend* abstraction offers interfaces for different kinds of backends such as caches, databases, and queues; Fig. 2 depicts a simple Cache interface. Unlike for services, a backend instantiation will have method implementations provided as part of its Blueprint compiler integration, e.g. memcached will provide a memcached client.

Blueprint imposes a *dependency injection* pattern on service implementations. A service can invoke the methods of another service or a backend (Line 12). However, a service is forbidden from instantiating other services or backends, which can only be received as constructor parameters (Line 8). Likewise, although invoking another service is simply a method call in the workflow spec (Line 12), the developer should not assume that other services and backends are running in the same address space, correspond to just a single instance, or are of a particular implementation. Services do not directly incorporate any scaffolding (e.g. configuring an RPC server) or bind specific instantiations (e.g. binding to

a memcached client library) as these are automatically integrated later by Blueprint’s compiler. Blueprint uses Go’s error-handling conventions to wrap and encapsulate errors that may be introduced from scaffolding, and Go’s context propagation for implementing scaffolding such as tracing.

The above restrictions are necessary for several reasons. First, the addressing scope of callee services can vary – they could be application-level instances in the same address space, or running in separate container instances in a different datacenter, requiring network calls and address translation. Second, scaffolding may interpose calls between services, e.g. to add functionality like tracing or to implement RPCs. Third, scaffolding may duplicate or replicate service instances in some way. In all cases, Blueprint’s compiler is responsible for later generating the code that instantiates, configures, and addresses services, at the application, process, and container granularities. §4.2 relates the above service abstraction to corresponding compiler abstractions.

Blueprint includes numerous out-of-the-box open-source applications, and once a developer has implemented an application’s workflow spec it only needs to be revisited if workflow logic needs to change. Changes to scaffolding and instantiations happen through Blueprint’s wiring spec.

Wiring Spec. The *wiring spec* declares the topology of the application, applies scaffolding, and configures instantiations. Fig. 3 depicts a wiring spec for three dependent DSB services [26]. Wiring is declared using a strongly-typed DSL with C-style macro support (Fig. 3). The wiring spec declares instances of services named in the workflow spec (Line 9) and links instance dependencies (Line 12); it also declares and links instantiations of backends (Line 7) and scaffolding (Line 2). The wiring spec also groups instances into deployment units such as processes and containers.

To enable scaffolding, the user refers to it using unique keywords and syntactic sugar in the wiring file (Line 2, 4). A corresponding compiler plug-in will be invoked during compilation to generate and modify code and other artifacts. Scaffolding can potentially apply to service instances, processes, or container images, depending on the specific feature it enables. For example, Zipkin tracing (Line 3) applies to service instances by wrapping with proxy classes.

A typical wiring spec is concise – tens of LoC – and easily modified by other Blueprint users. Users do not need to touch the workflow spec to enact changes in the wiring spec. Blueprint will recompile an application variant, potentially generating substantially modified code artifacts, without requiring manual intervention from the user.

Compiler Plugins. Scaffolding and instantiations are implemented as *compiler plugins*. Most applications will make use of Blueprint’s out-of-the-box compiler plugins; however a researcher wishing to prototype new functionality or improvements may wish to integrate that functionality with Blueprint by developing a compiler plugin. Compiler plugins

integrate with Blueprint in three places. First, a plugin can introduce keywords and syntactic sugar to the wiring spec. Second, as described in the next section, a plugin can extend Blueprint's IR to add new node types or extend existing node types. Third, a plugin provides logic for generating code, configuration, and other artifacts specific to the plugin. For example, Blueprint's gRPC plugin invokes the Protocol Buffers compiler and generates client and server wrapper classes. Blueprint is designed in a way that implementing a plugin is independent of the implementation of any other plugins. Most core concepts of Blueprint are implemented as compiler plugins, *e.g.* application-level Go service instances, Go processes, and Docker containers.

4.2 Intermediate Representation (IR)

The canonical representation of a Blueprint application is the compiler's Intermediate Representation (IR). Blueprint's compiler takes as input an application's workflow spec, wiring spec, and enabled compiler plug-ins. The IR of an application is a verbose and well-structured graph that represents the concrete layout and hierarchy of components along with their interactions. The IR of an application depends on both the workflow and wiring spec, and a change to just the wiring spec (*e.g.* to add replication) will result in a different IR. The IR is designed to support the flexibility and extensibility of the compiler. Fig. 4 depicts the IR graph for the wiring spec outlined in Fig. 3.

Component Nodes. Components are entities that will ultimately be instantiated in the generated system; they are represented as nodes in the IR. All services defined in the workflow spec have corresponding component nodes; likewise all backends and instantiations. Component nodes can exist at different granularities, *e.g.* representing an application-level service instance, a pre-defined binary (*e.g.* a MongoDB instance), or a pre-built container image (*e.g.* a ZipkinServer). IR nodes are typed and plugins may introduce new IR node types and implementations.

Namespace Nodes. Components of the same granularity can be grouped under a namespace node to create a component of coarser granularity. For example, a Go namespace node groups together application-level instances (*e.g.* service instances) into a single Go process. Similarly, a process namespace can be grouped into a container, and a container namespace into a deployment. During compilation, namespace nodes generate runnable artifacts (*e.g.* code, container images) that instantiate their contained components. Typing on nodes ensures that namespace nodes can only contain children of a compatible granularity. Blueprint can be extended with new namespaces; *e.g.* support for georeplication would introduce a region namespace; supporting C++ workflows would introduce a Cpp namespace.

Dependencies. Services in the workflow spec can invoke other services and components; in the IR these dependencies

are represented as edges between component nodes. RPC edges are directional indicating the caller-callee direction and declare the method signatures of the invocations.

Modifier Nodes. Scaffolding can interpose edges between components, *e.g.* to modify method signatures, add proxy functionality, or enable addressing across address space boundaries. We call these *modifier* nodes because they attach to component nodes and mutate the component's edges (*e.g.* to introduce client side and server side code). Modifiers must be opaque to the caller component whom expects a particular method signature from the callee; thus a modifier typically comprises a client-side transformation function and a corresponding server-side function that inverts the transformation (*e.g.* serialization and deserialization).

Visibility and Addressing. Dependent components can run in different processes, containers, or machines. For example, an application could be compiled as an all-in-one process, or using a container per service. Although there may be an edge between two components, it is possible that those components are not *visible* to each other, *e.g.* if a service has not been wrapped with an RPC server, it cannot receive remote invocations. Thus, edges between components are annotated with their visibility level. Nodes can expand the visibility of any edge traversing outside that node, *e.g.* an RPC modifier enables communication between processes.

Generators. A component declared in a wiring spec might correspond to a single concrete runtime instance (*e.g.* those in Fig. 3), or as a result of applying modifiers, to multiple runtime instances. For example, an autoscaling modifier might dynamically create and destroy multiple component instances at runtime. In general, *generator* nodes contain other nodes and represent instances that will be dynamically created at runtime. Generators restrict the visibility of contained nodes, since there will be multiple dynamically-generated instances of the contained nodes. Generators are typically coupled with functionality such as a load balancer to enable addressing of the dynamically created instances.

4.3 Compilation

Blueprint compiles an application in two steps. First, it processes the wiring spec and workflow spec to instantiate the specific IR graph representing the application and its topology. Second, it invokes compiler passes and scaffolding-specific plugins to generate the runnable artifacts. We explain both steps in detail below.

4.3.1 Wiring & Workflow Spec to IR

Blueprint parses the workflow spec to identify all workflow services that have been defined, and loads the definitions of standard library backends that can be instantiated. Next, Blueprint processes the wiring spec to extract the list of components instantiated in the wiring spec, creating the appropriate IR nodes for each. Blueprint applies modifiers

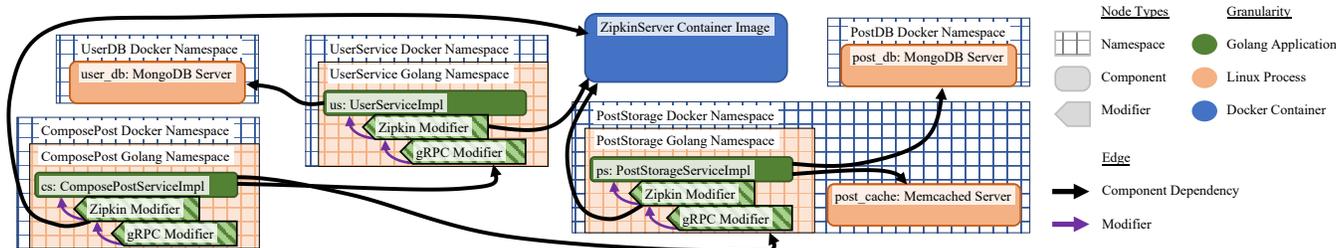


Fig. 4. A depiction of Blueprint’s IR for the three DSB services outlined in Fig. 3. Node shape is based on node type, and nodes are colored according to their granularity. Edges are color-coded according to the type of relationship.

to components by creating additional IR nodes representing the scaffolding. Blueprint then creates directed edges between components to encode the dependencies defined in the wiring spec. Blueprint then extracts the various component groupings and granularities to generate the namespace nodes as well as to add the visibility attributes to the dependency ages between the component nodes. Lastly, Blueprint performs a pass on the IR graph to allow modifier nodes to add, delete, or change nodes in the IR graph. For example, a replication modifier could duplicate the IR nodes representing a component, and insert a load balancer node.

4.3.2 IR to Implementation

Once the IR graph is constructed, Blueprint checks the visibility of edges, *i.e.* that any component that calls another component will be capable of doing so. Blueprint then proceeds to the artifact generation step. Each node of the IR graph will have plugin-specific logic for generating its own code, configuration, or artifacts needed for instantiating the component in the system. Blueprint traverses the IR graph, invoking plugins at IR nodes and collecting the artifacts that are generated.

Artifact Generation. Blueprint generates code and artifacts in a hierarchical manner, starting from the innermost nodes in the IR graph. For service nodes defined in the workflow spec, no extra code generation is required and only dependencies are gathered. For modifier nodes, the compiler invokes the plugin that corresponds to the node. The output of the previous node is passed as input to the plugin, allowing the plugin to potentially generate code that wraps, extends, or changes the previous output. For namespace nodes such as go processes or docker containers, generating code entails packaging code generated by the contained nodes, and generating code that instantiates those nodes.

As an example, consider the generation process for the ComposePostService in Fig. 3. The generation procedure starts at the ComposePostServiceImpl node, *cs*, in Fig. 4. This node simply gathers the code dependencies directly from the workflow spec where it is defined. The compiler then steps outwards to the ZipkinModifier which inspects the method signature list of *cs* and generates a wrapper class that adds

trace contexts to all methods. Next, the gRPC plugin generates protobuf RPC message definitions for the expanded *cs* methods, invokes the gRPC compiler, then generates client and server instantiation code. The compiler proceeds in inverse topological order and next invokes the Go Namespace plugin to package all contained code and generate a main method that instantiates the service, wrappers, RPC server, and clients to dependencies. The compiler proceeds similarly for process nodes and container nodes.

Resolving Dependencies. As part of the generation process, Blueprint gathers code dependencies across namespaces from the IR graph to ensure that remote components are addressable. For example, if a service invokes another over RPC, running in a different container, it must therefore include client code for the target service and its modifiers. Any node crossed by this edge must receive and forward the target service’s address as an argument (*i.e.* so that the target service binds to an address that the source service can dial, and so that docker containers publicly expose the relevant ports). If the remote components are not addressable by a service, Blueprint’s compiler will return an error citing that the edge between the two services lacks the necessary visibility.

5 Implementation

Blueprint is implemented in Go in 10,892 LoC, which includes Blueprint’s compiler, the wiring spec DSL, component interfaces and their implementations, debugging and logging features, and other features implemented as modifiers.

Blueprint’s compiler is implemented in 4062 LoC. Blueprint provides first-class support for Go workflow specs. We selected Go because it is designed specifically for high-performance RPC services, and has convenient mechanisms for handling concurrency, errors, and context propagation. Blueprint is not constrained to Go; the abstractions of Blueprint enable extension to multiple languages with no additional difficulty. Blueprint’s wiring spec is currently a Python-based DSL that also allows C-style macros; this is 771 LoC, and we are considering more flexible programmatic approaches for future work.

We reimplemented five applications from three microservice benchmark suites described in the literature: the SocialNetwork, Media, and HotelReservation applications from the DeathStar Benchmark (DSB) [26], TrainTicket [76], and the SockShop benchmark [47]. We present an analysis of the LoC effort required for porting these applications in §6.1. We additionally outline the features currently supported by Blueprint in §6.5 and the LoC of implementation required to implement the compiler plug-ins.

6 Evaluation

Our evaluation of Blueprint seeks to answer the following:

- Does Blueprint reduce effort for design space exploration (UC1)? (§6.1)
- Can Blueprint help reproduce emergent phenomena in microservice applications (UC2)? (§6.2)
- Does Blueprint reduce effort for prototyping improvements (UC3)? (§6.3)
- Are Blueprint-generated systems realistic? (§6.4)
- Is Blueprint easy to extend with new scaffolding and instantiations? (§6.5)
- What is the cost of Blueprint’s abstractions? (§6.6)

Experimental setup. All experiments use a cluster comprising eight machines, each with four Intel Xeon E7-8857V2 processors, 48 cores and 1.5 TB RAM. We deploy each service in a separate container. We use a simple open-loop workload generator that can be configured to exercise APIs of the generated system with a specified request rate and API distribution; this runs on a separate machine.

6.1 UC1: Design Space Exploration

Reducing Implementation Effort. To demonstrate that Blueprint makes it easy to implement realistic microservice systems not specifically designed for our evaluation, we have re-implemented five existing microservice applications in Blueprint. We selected these systems based on their popularity in microservice research as highlighted in §B.2. Tab. 1 shows the LoC needed to implement the workflow spec and wiring file for each application in Blueprint. We compare the LoC needed to those in the original implementations. Blueprint reduces the code footprint by 5–7× for each application by eliminating the need to implement scaffolding and instantiations alongside workflow code. In the original implementations, scaffolding was tightly coupled with workflow code, thus inflating the amount of code that a user needed to write. By decoupling the workflow specification from the scaffolding code and moving scaffolding code generation to the compiler, Blueprint reduces the volume of code required to implement a workflow. One beta user of Blueprint noted that this decoupling also made it easier for them to understand and implement the workflow specification.

System	Orig. (LoC)	Blueprint (LoC)		Reduction
		Spec	Wiring	
DSB SocialNetwork	8 209	1 478	57	5.4×
DSB Media	7 794	1 401	42	5.4×
DSB HotelReservation	5 160	679	63	7.0×
TrainTicket	54 466	9 639	166	5.6×
SockShop	13 987	2 261	40	6.1×

Tab. 1. Lines of code of original and Blueprint implementations of popular open-source microservice systems

Changing workflow specs. We compare the LoC required to make a change to the design of the application in the original system and compare that to the effort to implement the same change in the Blueprint implementation of the application. In pull request #101 of DSB SocialNetwork [19], the authors inverted caller-callee relationships between ComposePostService and TextService, UniqueIDService, UserService, and MediaService to improve application performance. They modified 5,140 LoC. We implemented the same change in the Blueprint version of the system by modifying 288 LoC of workflow spec, and 7 LoC of wiring spec — a 17× reduction. The substantial difference in code changes arises due to Blueprint’s separation of concerns: in the original implementation, changing interfaces between services required changing scaffolding and instantiation code, which was tightly coupled with workflow code. However, with Blueprint, scaffolding and instantiation code changes are handled by the compiler, and only the workflow specification required manual changes.

Changing scaffolding and instantiations. Blueprint makes it easier to enable or disable new scaffolding in an application. Based on our survey in §B.2, we found that a popular change in existing microservice systems is to enable or disable tracing [5, 13, 32, 37]. For example, disabling tracing in a variant of DSB SocialNetwork required 418 LoC [37]. In contrast, the same change required 5 LoC wiring spec change for the Blueprint implementation of DSB SocialNetwork. This small wiring spec change automatically removes ~2KLoC from the generated system, including all tracing source code modifications and configuration files needed to enable tracing in the runtime.

Performance-Driven Design Exploration. Finally we perform a study requiring many configure, build, and deploy iterations. We first study the performance impact of different choices in DSB HotelReservation and DSB SocialNetwork [26]. Fig. 5 shows the performance impact of changes along two dimensions: (i) the RPC framework (gRPC or Thrift); and (ii) the size of the clientpool (relevant only for Thrift, as gRPC multiplexes connections on a single connection). We find that gRPC outperforms Thrift for both applications and find marginal differences when varying the clientpool size.

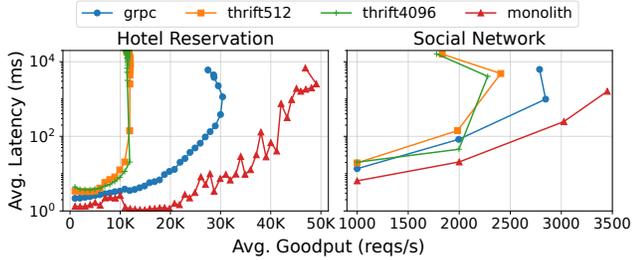


Fig. 5. Blueprint enables easy performance-driven design exploration.

Next, we use Blueprint to generate monolithic versions of both applications, where all services run in a single process and communicate directly through function calls. This allows us to quantify the performance cost of breaking the application down into a microservice architecture. In both cases, we run all services on a single machine. The monolith line in Fig. 5 shows that the monolith version outperforms the microservice version of the application. This enables an empirical decision for when the complexity of a microservice architecture is justified from a performance point of view.

To generate these variants, we only needed to modify 5–10 LoC in the wiring spec, illustrating how Blueprint enables design space exploration with minimal manual effort.

6.2 UC2: Eliciting Emergent Phenomena

Through two case-studies we demonstrate that Blueprint is capable of generating systems for exploring emergent phenomena, namely *metastability failures* [33] and *cross-system inconsistency* [61]. We modify the Blueprint implementations of DSB HotelReservation and DSB SocialNetwork to exhibit these specific emergent phenomena; for readability we refer to these simply as HotelReservation and SocialNetwork.

6.2.1 Case Study I: Metastability Failures

We adapt HotelReservation and SocialNetwork to showcase the four different kinds of metastability failures [33]. The required changes described below are to the wiring spec and amount to at most 3 LoC per failure type.

Load spike trigger workload amplification (Type 1).

We modify HotelReservation to add a 500 ms timeout to every inter-service RPC. We also modify the services to retry up to 10 times on error. We start with a 10 K requests/s (Rps) workload for 60 s, then increase to 30 KRps for 30 s, and then revert to 10 KRps. Fig. 6a shows the mean and 99th percentile service latencies over time. At the 1-minute mark, the sudden workload increase triggers the majority of requests to time out, in turn causing more requests to be generated due to retries. This trigger keeps the system in a metastable state and even after decreasing the load, the system does not recover to a stable state.

Load spike trigger capacity degradation amplification (Type 2).

To induce this type of metastability failure, the authors [33] limited the maximum service heap size. We experiment with HotelReservation’s ReservationService. As Go offers no direct control over heap size, we instead increase the garbage collection (GC) frequency by causing the GC to trigger whenever the heap is 75% full instead of the default 100% (for this, we set the environment variable GOGC to 75). We run a 20 KRps workload for 10 mins; at the 5 min mark we introduce low-level CPU contention for 30 s using FIRM’s anomaly injector [55]. Fig. 6b shows mean service latency over time. Here, the CPU contention acts as a trigger by increasing the GC duration, which results in frequent stop-the-world GC phases, causing other requests to start timing out and generate more retries. This metastable state also persists after the CPU contention disappears.

Capacity decreasing trigger workload amplification (Type 3).

To induce this metastability failure, we first modify HotelReservation to have 1 s timeouts and up to 10 retries on every RPC. We run a 24 KRps workload for 2 mins. After 1 min, we inject low-level resource contention with FIRM’s [55] anomaly injector for 30 s to decrease available CPU capacity. Fig. 6c shows the mean and 99th percentile service latencies over time. CPU contention starting at 1 min causes timeouts leading to retries that overload the system and keep it in a metastable state after CPU contention disappears.

Fig. 7 illustrates vulnerability for different request rates, trigger durations, and maximum retries. At higher request rates, even a short trigger can cause the system to move into a metastable failure state. In contrast, at lower request rates, short triggers only cause transient issues. Fewer retries only minimally increase the tolerable trigger duration.

Capacity decreasing trigger capacity degradation amplification (Type 4).

We modify SocialNetwork with an internal 1 s timeout and up to 10 retries. We run this experiment in two phases. First, we fill the UserTimelineCache with all content from the userTimelineDatabase. In the second phase, we run a 3 KRps workload for 2 mins. At 1 min, we flush the UserTimelineCache which then causes future requests to query the database. Fig. 6d again shows the mean and 99th percentile service latencies along with the observed cache miss rate. The cache flush at 1 min overloads the database and causes cascading timeouts and retries. This overload prevents the cache from repopulating and keeps the system in a metastable failure state.

In all cases, Blueprint enables rapid exploration of different designs, such as adding timeouts and retries, which, in turn, enables the reproduction and analysis of metastable failures.

6.2.2 Case Study II: Cross-System Inconsistency

We add replication scaffolding to SocialNetwork to elicit cross-system inconsistencies. Cross-system inconsistencies

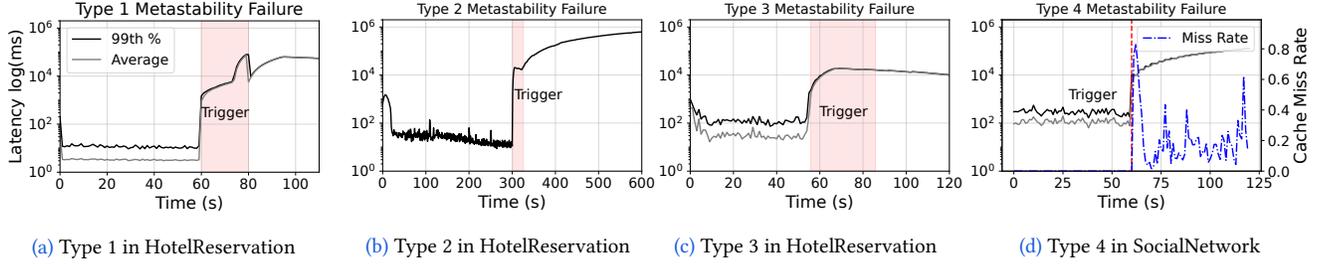


Fig. 6. Four metastability failure types [33] demonstrated in the Blueprint versions of DSB HotelReservation and DSB SocialNetwork.

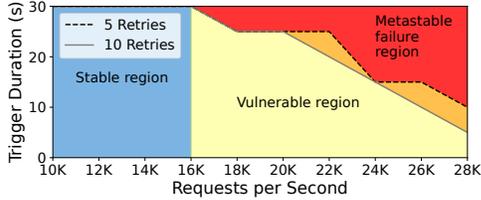


Fig. 7. Metastability Vulnerability Analysis for HotelReservation.

occur when delays in synchronization of replicated data stores such as databases and caches cause read requests for an object to return incorrect results. By default SocialNetwork has no replication, so we add replication to userTimeline-Database and UserTimelineService, and modify the GatewayService to use the replicated version of the service. These changes only touch 4 LoC in the wiring spec.

We compare the replication-enabled SocialNetwork to the initial Blueprint SocialNetwork. We use a 100 Rps workload consisting of 100 % ComposePost requests. For each successful request, we read the user timeline of the post creator after a wait time ranging between 0 s and 1 s at 100 ms intervals. A response without the new post is a cross-system inconsistency. Fig. 8 shows the measured fraction of inconsistencies with increasing wait times for the original SocialNetwork and the replicated version. As expected, the non-replicated version always reads consistent results, whereas the replicated version incurs a small fraction of inconsistencies [12].

Overall, this case study demonstrates how Blueprint applications are amenable to change and enable users to modify existing applications to reproduce emergent phenomena with minimal effort.

6.3 UC3: Prototyping Improvements

In this section we evaluate how amenable Blueprint is for supporting prototyping and evaluation of improvements in two ways: (i) reproducing the prototyping required for a solution performed in existing research; and (ii) prototyping a new solution for an emergent phenomenon.

Reproducible Research. To understand how Blueprint can aid researchers in making changes, we select a mutation that was manually added by researchers to an existing

microservice application, and reproduce that mutation in the Blueprint implementation of the application. Concretely, Sifter [40] manually mutates the DSB Social Network application to add X-Trace [23]. X-Trace is a distributed tracing framework, but it does not conform to OpenTelemetry APIs and cannot reuse the existing Jaeger instrumentation of DSB SocialNetwork. Consequently, the authors of Sifter manually extended DSB SocialNetwork to add X-Trace support. This comprised 1,289 LoC changed over a 2 week period, based on commit timestamps.

We implemented the same change in Blueprint, which required (1) extending Blueprint’s compiler to support X-Trace (a 1-time task); and (2) enabling X-Trace for the Blueprint SocialNetwork application. The latter required 3 LoC change to the wiring spec of the SocialNetwork application. This reduction occurred because the vast majority of code to support X-Trace is templatable scaffolding code which can be easily incorporated in the compiler. Implementing X-Trace within Blueprint’s compiler required 433 LoC.

To evaluate if Blueprint’s modifications to systems yield the same results as the modifications to original systems, we contacted the authors of Sifter to obtain their experiment code and reproduced Figure 6 from the Sifter paper [40] using the Blueprint-generated SocialNetwork application. In the original experiment, the authors recorded the loss and sampling probability for a sequence of 1000 ComposePost API requests, and at five separate instances they had inserted anomalous requests. Similarly, in our experiment we generated anomalous requests with the Blueprint generated DSB-SN and repeated the above experiment. In Fig. 9, the spikes of high probability of selection correspond to the anomalous requests. This shows that Blueprint generated systems can reproduce the same results as the original systems.

Prototyping New Solutions. In this experiment we demonstrate how Blueprint can help in prototyping and integrating novel solutions in applications. We particularly focus on the Type 1 Metastable Failure first introduced in §6.2.

To address the Type 1 Metastable Failure, we prototype a solution based on *circuit-breakers*. A circuit-breaker prevents new requests from being sent if the moving-average failure rate of requests is higher than a specified threshold. We

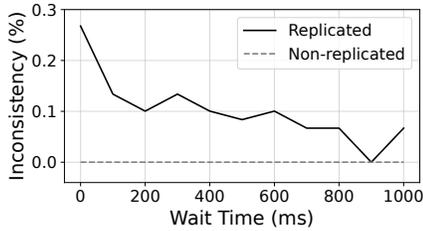


Fig. 8. Cross-system inconsistencies arise in SocialNetwork when enabling replication.

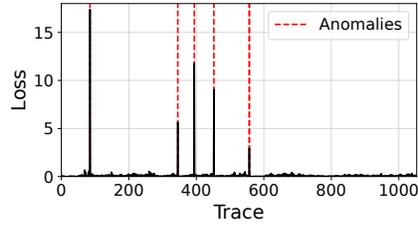


Fig. 9. Blueprint’s reproduction of Fig. 6 from the Sifter paper [40] in SocialNetwork.

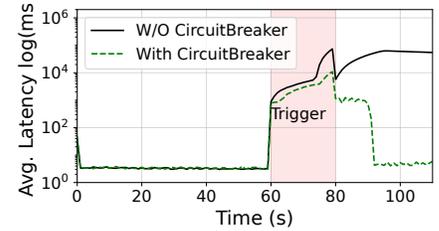


Fig. 10. Prototype Solution for Type 1 Metastability failure in HotelReservation.

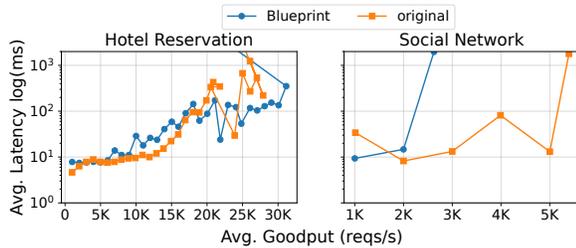


Fig. 11. Performance comparison of Blueprint systems with their original counterparts

implement a circuit-breaker feature and introduce it to Blueprint’s compiler as a new type of scaffolding. This required 126 LoC. Enabling circuit breakers for the HotelReservation application required only 2 LoC change to its wiring spec.

Fig. 10 shows how adding circuit breakers can potentially prevent the system from going into a metastable failure state. The circuit-breaker-enabled version of the application experiences the same latency increase at $t = 60$, but due to the circuit-breaker triggering, the application is able to avoid entering a metastable failure state and returns to normal at $t \approx 90$. Overall, this example demonstrates how Blueprint can be useful in prototyping novel solutions for phenomena.

6.4 Generating Realistic Systems

We now evaluate Blueprint’s ability to generate real systems useful for meaningful evaluation. For this we measure and compare the performance of two of the Blueprint-generated systems applications above to that of the original systems.

HotelReservation. We compare the performance of the Blueprint implementation of HotelReservation to the original DSB implementation. The original DSB HotelReservation application is implemented in Go, enabling a direct comparison between it and the Blueprint-generated application. To exercise the systems, we run a mixed workload of 60% hotels, 38% recommendations, 1% user, and 1% recommend requests. We run the workload for different request rates ranging from 1 Krps to 30 KRps for a duration of 1 min each. Fig. 11 shows the latency-throughput profile of both systems. The Blueprint generated system has a similar performance to the original manually implemented system.

Backend	Interface	Compiler
Cache	12	0
NoSQLDB	27	0
RelDB	22	0
Queue	12	0
Tracer	45	0
Deployer	3	46
RPC	11	152
HTTP	11	146

Tab. 2. Lines of Code required for adding a backend interface.

SocialNetwork. We also compare the performance of the Blueprint implementation of SocialNetwork to the original DSB implementation. The original DSB implementation uses a mix of C++ and Lua with an nginx gateway web server whereas the Blueprint implementation uses Go’s default HTTP web server. To exercise both systems, we run a mixed workload of 60% ReadHomeTimeline, 30% ReadUserTimeline, and 10% ComposePost requests. We run the workload for different request rates ranging from 1 KRps to 6 KRps for a duration of 1 min each. Fig. 11 shows the latency-throughput profile for both systems under the aforementioned workload. In this scenario, the original system outperforms the Blueprint generated system. We attribute this to two compounding factors. First, the original system is implemented in C++ whereas the Blueprint version is in Go. Go incurs garbage collection overhead and relies on different libraries for many core microservices building blocks. Second, the Blueprint implementation does not use any Redis-specific specialized array operations. This illustrates a drawback of Blueprint– it requires interacting with backends through common interfaces.

Overall, these results illustrate that Blueprint can generate microservice systems whose performance compares closely to that of handwritten systems.

6.5 Extensibility of Blueprint

Adding backends and instantiations. Tab. 2 shows the LoC in the interface of various Blueprint backends and

Type	Instantiation	Impl	Compiler
Cache	Redis	76	140
Cache	Memcached	76	142
NoSQLDB	MongoDB	288	140
RelDB	MySQL	91	140
Queue	RabbitMQ	50	111
Tracer	Jaeger	28	145
Tracer	Zipkin	28	145
Deployer	Docker	74	0
Deployer	Kubernetes	45	0
Deployer	Ansible	439	0
RPC	GRPC	673	0
RPC	Thrift	636	0
HTTP	Go's net/http	271	0

Tab. 3. Lines of Code required for adding a new instantiation.

Plugin	Compiler (LoC)	Stdlib (LoC)
Retry	123	0
Tracing	284	45
p-Replication	52	0
ClientPools	145	55
X-Trace	364	69
CircuitBreaker	126	0
LoadBalancer	208	19

Tab. 4. Lines of code required for adding a new compiler plugin.

Tab. 3 shows the LoC for the instantiations currently available for each backend. Adding a new backend generally requires <100 LoC for implementing the lightweight interface. Adding an instantiation also requires a small amount of code (<200 LoC) in the compiler. Each of these is a one-time effort, that can be leveraged by subsequent Blueprint applications.

Instantiations of certain backends require more code than others. For example, instantiations of the RPC and backends require more code than average as these instantiations must correctly generate the code for communicating over the network, establishing connections, and running servers.

Adding new plugins. Blueprint’s modifier abstraction allows developers to add new scaffolding. Tab. 4 shows the plugins currently available in Blueprint and the LoC in their implementations. Some plugins require changes only in the Blueprint toolchain whereas others also require an additional runtime library component. The LoC vary across features from 46 to around 440.

Adding new plugins is harder than directly implementing the scaffolding hard-coded within an application. Nonetheless, this addition is a one-time cost that amortizes the effort of reimplementing scaffolding in all later systems.

6.6 Cost of Blueprint

Compilation time. Tab. 5 shows the time taken by Blueprint to generate systems. Blueprint can generate small to

System Name	Gen Time(s)	Num Services
DSB-SN	1.172	28
DSB-MM	1.698	33
DSB-HR	1.281	18
TrainTicket	3.723	67
SockShop	0.925	14
Alibaba-TraceSet	707	2882

Tab. 5. Time taken by Blueprint for generating the system including services, caches, databases, queues, tracers

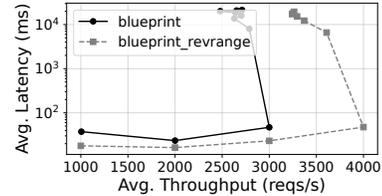


Fig. 12. DSB-SN cache choice exploration

medium sized system within seconds. As there are no existing large open-source microservice systems, we generated a large-scale microservice application using the Alibaba service topology in the Alibaba trace dataset [44]. For this, we omitted the caches and databases and only focused on stateless services. In total, the resulting workflow and wiring spec had 2.8K service instances. To generate a variant implementation, Blueprint took around 12 minutes. Overall, the compilation time is proportional to the number of service instances in the wiring spec and the density of the service topology. These results demonstrate that Blueprint enables researchers and practitioners to quickly (re-)compile applications, thus supporting rapid Configure, Build, and Deploy cycles. These results may be further improved through compiler optimizations and incremental compilation.

Cost of abstractions. For each type of backend supported by Blueprint, services interact with the backend through a generalized interface. The interface is selected such that backend instances can be opaquely reconfigured. Yet many backends do provide broader interfaces with specialized operations that are more efficient for certain workloads. Blueprint’s approach discourages services from using such operations, in the interest of reconfiguration.

To demonstrate the impact of using more general but less efficient APIs, we implement an extended Cache interface that provides access to Redis’ specialized cache operations. We use this extended interface in the SocialNetwork application, we execute a 100% ReadHomeTimeline workload for 1 minute for request rates ranging from 1 KRps to 6 KRps. With the extended interface, the application observes a 33% increase in throughput as shown in Fig. 12.

7 Discussion

Debugging. Debugging generated code and the workflow specification can be a challenging task. To aid developers in debugging workflow specification code, Blueprint provides default implementations of the various components called *null implementations*. These implementations provide a basic interface against which the core application specification can be tested without worrying about the deployment of the system. These implementations are attached to the workflow specification in the wiring spec and can be iteratively replaced with the real choices once the developer is confident in the correctness of their application code.

Language Heterogeneity. Usually, microservice systems contain services implemented in different languages. Currently, Blueprint generates services only in Go. Generating services in other languages is purely an implementation challenge and we believe that the current design can be extended to generate code in other languages. This would require compiler plugins to support generation of scaffolding in multiple languages in order to correctly mix code of different languages in an application.

Generating Production Systems. Blueprint targets microservice experimentation and prototyping use cases, rather than generating production-ready microservice applications. It remains an open question whether Blueprint is a suitable toolchain for developing production-ready microservice applications. Currently, Blueprint’s approach stands at odds with the microservices architectural approach: microservices are typically developed by highly distributed teams operating independently; yet Blueprint imposes a centralized configuration and deployment step through its wiring spec. We are currently exploring approaches to decomposing and distributing the wiring spec. A further concern is the cost of Blueprint’s abstractions for backends, which might be too high a price to pay for production systems. However, we believe that Blueprint could be used in production to quickly home in on a concrete set of choices that satisfies the developer’s requirements. The Blueprint-generated system could then further be fine-tuned manually to obtain a production-grade system.

8 Related Work

Microservice benchmark suites have gained considerable popularity in recent years [26, 66, 71, 76]. Each system corresponds to only one concrete point in the design space and as they were not designed to be configurable, they are inadequate for tasks that would require exploring the design space of microservices.

Several existing tools support generation of microservice systems by providing a DSL or some other programming language [4, 21, 27, 38, 49, 53, 56–59, 62–64, 67, 68, 73].

However, the tools are designed for one-shot generation of microservices and select a single dimension to provide flexibility. The most prominent example is Google’s ServiceWeaver [27], which provides flexibility along the deployment dimension allowing users of the tool to deploy the same application as a suite of microservices or as a monolithic application. However, like the other tools, ServiceWeaver is a poor fit CBD use cases that require modifying the application along dimensions other than the deployment dimension.

Some existing tools, such as SpringBoot [65] or Dapr [1] aid in developing microservice systems by separating out reusable infrastructure components from the core implementation of the application, allowing users to select infrastructure building blocks that can be applied to an application. However, the SpringBoot framework makes binding choices along the deployment and communication dimensions of the microservice design space making SpringBoot a poor fit for CBD use cases. Unlike the SpringBoot framework, Dapr allows users to switch between the deployment and communication dimensions but the user must change them manually resulting in high manual effort.

Two of the key features that enables Blueprint to quickly reconfigure applications are the notions of reusability and dependencies. First, the idea of reusability has been inspired from the The Flux OSKit [24]. Similar to Flux OSKit’s suite of reusable OS components, Blueprint also uses commonly used microservice libraries and components as reusable building blocks for microservice applications. In contrast, Blueprint does not require glue code to be handwritten for each component for each new application, instead relying on its compiler abstractions and code generation to correctly handle composition of components. Second, Blueprint uses Dependency Injection [54] to correctly generate applications by not allowing instantiation of dependent components in the workflow specification of applications. The dependencies of any given service in a Blueprint application are generated at compilation time.

9 Conclusion

We have introduced Blueprint, a toolchain for developing highly reconfigurable microservice applications. We have demonstrated Blueprint for several use cases that involve configuring, building, and deploying variants of microservice applications with modified designs. Compared to existing benchmark applications, Blueprint substantially eases development and reconfiguration by providing a strong separation of concerns between an application’s workflow, scaffolding infrastructure, and implementation choices. Blueprint is open-source, and we hope that its adoption will make it easier for future work to understand and improve the performance and correctness guarantees of microservice systems. Blueprint is available at <https://gitlab.mpi-sw.s.org/cld/blueprint/blueprint-compiler>.

Acknowledgements

We would like to thank our shepherd, Rebecca Isaacs and our anonymous reviewers for helping us shape up the final version of the paper. We would also like to thank and acknowledge the efforts of the anonymous evaluators of the Artifact Evaluation Committee. We would like to thank Gerd Alliu, for helping with an initial Blueprint versions of the TrainTicket and SockShop applications. We would also like to thank Rodrigo Fonseca, Roberta De Viti, and Matheus Stolet for providing us with feedback.

References

- [1] Dapr: Distributed application runtime. <https://dapr.io/>.
- [2] AcmeAir. Acmeair. <https://github.com/acmeair>.
- [3] P. Ajoux, N. Bronson, S. Kumar, W. Lloyd, and K. Veeraraghavan. Challenges to adopting stronger consistency at scale. In *15th Workshop on Hot Topics in Operating Systems (HotOS XV)*, 2015.
- [4] I. K. Aksakalli, T. Celik, A. B. Can, and B. Tekinerdogan. A model-driven architecture for automated deployment of microservices. *Applied Sciences*, 11(20):9617, 2021.
- [5] V. Anand and J. Mace. Deathstarbench - modified with x-trace. <https://github.com/JonathanMace/DeathStarBench/>.
- [6] asc lab. Asclab micronaut poc - lab insurance sales portal. <https://github.com/asc-lab/micronaut-microservices-poc/>.
- [7] Azure. Retry storm antipattern. Retrieved September 2022 from <https://learn.microsoft.com/en-us/azure/architecture/antipatterns/retry-storm/>, 2021.
- [8] Azure. Chatty i/o. Retrieved September 2022 from <https://learn.microsoft.com/en-us/azure/architecture/antipatterns/chatty-io/>, 2022.
- [9] Azure. No caching antipattern. Retrieved September 2022 from <https://learn.microsoft.com/en-us/azure/architecture/antipatterns/no-caching/>, 2022.
- [10] Azure. Noisy neighbour. Retrieved September 2022 from <https://learn.microsoft.com/en-us/azure/architecture/antipatterns/noisy-neighbor/noisy-neighbor>, 2022.
- [11] Azure. Technology choices for azure solutions. Retrieved September 2022 from <https://learn.microsoft.com/en-us/azure/architecture/guide/technology-choices/technology-choices-overview>, 2022.
- [12] P. Bailis, S. Venkataraman, M. J. Franklin, J. M. Hellerstein, and I. Stoica. Probabilistically bounded staleness for practical partial quorums. *Proceedings of the VLDB Endowment*, 5(8):776–787, 2012.
- [13] Barber0. Deathstarbench no tracing. Retrieved August 2022 from https://github.com/Barber0/DeathStarBench/tree/no_tracing, 2021.
- [14] R. Blum and R. Singh. Data integrity: What you read is what you wrote. Retrieved September 2022 from <https://sre.google/sre-book/data-integrity/>.
- [15] N. Bronson, A. Aghayev, A. Charapko, and T. Zhu. Metastable failures in distributed systems. In *Proceedings of the Workshop on Hot Topics in Operating Systems*, pages 221–227, 2011.
- [16] A. Cockcroft. The evolution of microservices. (April 2016). Retrieved October 2020 from <https://www.slideshare.net/adriancockcroft/evolution-of-microservices-craft-conference>, 2016.
- [17] A. Cockcroft. Microservices workshop: Why, what, and how to get there. (April 2016). Retrieved October 2020 from <https://www.slideshare.net/adriancockcroft/microservices-workshop-craft-conference>, 2016.
- [18] J. Dean and L. A. Barroso. The tail at scale. *Communications of the ACM*, 56(2):74–80, 2013.
- [19] delimitrou. Deathstarbench social network application refactor for better performance. Retrieved August 2022 from <https://github.com/delimitrou/DeathStarBench/pull/101>, 2021.
- [20] M. Dinga. An empirical evaluation of the energy and performance overhead of monitoring tools on docker-based systems. Retrieved August 2022 from <https://github.com/MadalinaDinga/thesis-2022-monitoring-tools-integration-with-train-ticket-replication-package/>, 2022.
- [21] T. F. Düllmann and A. Van Hoorn. Model-driven generation of microservice architectures for benchmarking performance and resilience engineering approaches. In *Proceedings of the 8th acm/spec on international conference on performance engineering companion*, pages 171–172, 2017.
- [22] J. Faro. Deathstarbench metastability failure fork. Retrieved August 2022 from <https://github.com/jfaro/DeathStarBench>, 2022.
- [23] R. Fonseca, G. Porter, R. H. Katz, and S. Shenker. X-trace: A pervasive network tracing framework. In *4th {USENIX} Symposium on Networked Systems Design & Implementation ({NSDI} 07)*, 2007.
- [24] B. Ford, G. Back, G. Benson, J. Lepreau, A. Lin, and O. Shivers. The flux oskit: A substrate for kernel and language research. In *Proceedings of the sixteenth ACM symposium on Operating systems principles*, pages 38–51, 1997.
- [25] Y. Gan, S. Dev, D. Lo, and C. Delimitrou. Sage: Leveraging ML To Diagnose Unpredictable Performance in Cloud Microservices. In *Workshop on ML for Computer Architecture and Systems (MLArchSys)*, June 2020.
- [26] Y. Gan, Y. Zhang, D. Cheng, A. Shetty, P. Rathi, N. Kataraki, A. Bruno, J. Hu, B. Ritchken, B. Jackson, et al. An open-source benchmark suite for microservices and their hardware-software implications for cloud & edge systems. In *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems*, pages 3–18, 2019.
- [27] S. Ghemawat, R. Grandl, S. Petrovic, M. Whittaker, P. Patel, I. Posva, and A. Vahdat. Towards modern development of cloud applications. In *Proceedings of the 19th Workshop on Hot Topics in Operating Systems*, pages 110–117, 2023.
- [28] GoogleCloudPlatform. Online boutique, fka hipstershop. <https://github.com/GoogleCloudPlatform/microservices-demo>.
- [29] E. Haddad. Service-oriented architecture: Scaling the uber engineering codebase as we grow. (September 2015). Retrieved October 2020 from <https://eng.uber.com/service-oriented-architecture/>, 2015.
- [30] M. Hashemi. (January 2017). Retrieved February 2021 from https://blog.twitter.com/engineering/en_us/topics/infrastructure/2017/the-infrastructure-behind-twitter-scale.html, 2017.
- [31] G. Heiser. System benchmarking crimes. <https://www.cse.unsw.edu.au/~gernot/benchmarking-crimes.html>.
- [32] Hilbert-Yaa. Deathstarbench local tracing. Retrieved August 2022 from <https://github.com/Hilbert-Yaa/DeathStarBench/tree/proto-tracing>, 2021.
- [33] L. Huang, M. Magnusson, A. B. Muralikrishna, S. Estyak, R. Isaacs, A. Aghayev, T. Zhu, and A. Charapko. Metastable failures in the wild. In *16th USENIX Symposium on Operating Systems Design and Implementation (OSDI 22)*, pages 73–90, 2022.
- [34] L. Huang and T. Zhu. tprof: Performance profiling via structural aggregation and automated analysis of distributed systems traces. In *Proceedings of the ACM Symposium on Cloud Computing*, pages 76–91, 2021.
- [35] D. Huye, Y. Shkuro, and R. R. Sambasivan. Lifting the veil on {Meta’s} microservice architecture: Analyses of topology and request workflows. In *2023 USENIX Annual Technical Conference (USENIX ATC 23)*, pages 419–432, 2023.
- [36] istio. Bookinfo application. <https://istio.io/latest/docs/examples/bookinfo/>.
- [37] ivanium. Deathstarbench social network no tracing. Retrieved August 2022 from <https://github.com/ivanium/DeathStarBench/commit/01360df6653e12982335838c877cc4178a518987>, 2021.
- [38] jhipster. Jhipster. Retrieved August 2022 from <https://github.com/jhipster/generator-jhipster>, 2019.

- [39] A. Kalia, M. Kaminsky, and D. Andersen. Datacenter {RPCs} can be general and fast. In *16th USENIX Symposium on Networked Systems Design and Implementation (NSDI 19)*, pages 1–16, 2019.
- [40] P. Las-Casas, G. Papakerashvili, V. Anand, and J. Mace. Sifter: Scalable sampling for distributed traces, without feature engineering. In *Proceedings of the ACM Symposium on Cloud Computing*, pages 312–324, 2019.
- [41] A. Lesniak, R. Laigner, and Y. Zhou. Enforcing consistency in microservice architectures through event-based constraints. In *Proceedings of the 15th ACM International Conference on Distributed and Event-based Systems*, pages 180–183, 2021.
- [42] J. Loff. Deathstarbench inconsistency. Retrieved August 2022 from <https://github.com/jfloff/antipode-deathstarbench>, 2022.
- [43] J. Loff, D. Porto, C. Baquero, J. Garcia, N. Preguiça, and R. Rodrigues. Transparent cross-system consistency. In *Proceedings of the 3rd International Workshop on Principles and Practice of Consistency for Distributed Data*, pages 1–4, 2017.
- [44] S. Luo, H. Xu, C. Lu, K. Ye, G. Xu, L. Zhang, Y. Ding, J. He, and C. Xu. Characterizing microservice dependency and performance: Alibaba trace analysis. In *Proceedings of the ACM Symposium on Cloud Computing*, pages 412–426, 2021.
- [45] J. Mace, P. Bodik, R. Fonseca, and M. Musuvathi. Retro: Targeted resource management in multi-tenant distributed systems. In *12th USENIX Symposium on Networked Systems Design and Implementation (NSDI 15)*, pages 589–603, 2015.
- [46] A. Mcdade. Significant outage for amazon web services stalls netflix, delta airlines, others. Retrieved September 2022 from <https://www.newsweek.com/significant-outage-amazon-web-services-stalls-netflix-delta-airlines-others-1657077>, 2021.
- [47] microservices demo. Retrieved August 2022 from <https://github.com/microservices-demo/microservices-demo>, 2016.
- [48] J. C. Mogul. Emergent (mis) behavior vs. complex software systems. *ACM SIGOPS Operating Systems Review*, 40(4):293–304, 2006.
- [49] F. Montesi, C. Guidi, and G. Zavattaro. Service-oriented programming with jolie. In *Web Services Foundations*, pages 81–107. Springer, 2014.
- [50] S. Needleman. Amazon outage disrupts lives, surprising people about their cloud dependency. Retrieved September 2022 from <https://www.wsj.com/articles/amazon-outage-disrupts-lives-surprising-people-about-their-cloud-dependency-11638972001>, 2021.
- [51] L. Nolan. Why health checks are like sidewalks. (December 2022). Retrieved April 2023 from <https://www.usenix.org/publications/loginonline/why-health-check-sidewalk>, 2022.
- [52] J. Park, B. Choi, C. Lee, and D. Han. Graf: a graph neural network based proactive resource allocation framework for slo-oriented microservices. In *Proceedings of the 17th International Conference on emerging Networking EXperiments and Technologies*, pages 154–167, 2021.
- [53] K. Perera and I. Perera. A rule-based system for automated generation of serverless-microservices architecture. In *2018 IEEE International Systems Engineering Symposium (ISSE)*, pages 1–8. IEEE, 2018.
- [54] D. R. Prasanna. Dependency injection. 2009.
- [55] H. Qiu, S. S. Banerjee, S. Jha, Z. T. Kalbarczyk, and R. K. Iyer. {FIRM}: An intelligent fine-grained resource management framework for slo-oriented microservices. In *14th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 20)*, pages 805–825, 2020.
- [56] F. Rademacher, S. Sachweh, and A. Zündorf. Aspect-oriented modeling of technology heterogeneity in microservice architecture. In *2019 IEEE International conference on software architecture (ICSA)*, pages 21–30. IEEE, 2019.
- [57] F. Rademacher, J. Sorgalla, S. Sachweh, and A. Zündorf. Towards a viewpoint-specific metamodel for model-driven development of microservice architecture. *arXiv preprint arXiv:1804.09948*, 2018.
- [58] F. Rademacher, J. Sorgalla, S. Sachweh, and A. Zündorf. Specific model-driven microservice development with interlinked modeling languages. In *2019 IEEE International Conference on Service-Oriented System Engineering (SOSE)*, pages 57–5709. IEEE, 2019.
- [59] F. Rademacher, J. Sorgalla, P. N. Wizenty, S. Sachweh, and A. Zündorf. Microservice architecture and model-driven development: Yet singles, soon married (?). In *Proceedings of the 19th International Conference on Agile Software Development: Companion*, pages 1–5, 2018.
- [60] M. R. S. Sedghpour, C. Klein, and J. Tordsson. Service mesh circuit breaker: From panic button to performance management tool. In *Proceedings of the 1st Workshop on High Availability and Observability of Cloud Systems*, pages 4–10, 2021.
- [61] X. Shi, S. Pruett, K. Doherty, J. Han, D. Petrov, J. Carrig, J. Hugg, and N. Bronson. {FlightTracker}: Consistency across {Read-Optimized} online stores at facebook. In *14th USENIX Symposium on Operating Systems Design and Implementation (OSDI 20)*, pages 407–423, 2020.
- [62] J. Sorgalla. Ajil: A graphical modeling language for the development of microservice architectures. In *Extended Abstracts of the Microservices 2017 Conference*, 2017.
- [63] J. Sorgalla, F. Rademacher, S. Sachweh, and A. Zündorf. Model-driven development of microservice architecture: An experiment on the quality in use of a uml-and a dsl-based approach. 2020.
- [64] J. Sorgalla, P. Wizenty, F. Rademacher, S. Sachweh, and A. Zündorf. Applying model-driven engineering to stimulate the adoption of devops processes in small and medium-sized development organizations. *SN Computer Science*, 2(6):1–25, 2021.
- [65] spring.io. Spring boot. Retrieved August 2022 from <https://spring.io/projects/spring-boot>.
- [66] A. Sriraman and T. F. Wenisch. μ suite: A benchmark suite for microservices. In *2018 IEEE International Symposium on Workload Characterization (IISWC)*, pages 1–12. IEEE, 2018.
- [67] A. Suljkanović, B. Milosavljević, V. Indjić, and I. Dejanović. Developing microservice-based applications using the silvera domain-specific language. *Applied Sciences*, 12(13):6679, 2022.
- [68] B. Terzić, V. Dimitrieski, S. Kordić, G. Milosavljević, and I. Luković. Development and evaluation of microbuilder: a model-driven tool for the specification of real microservice software architectures. *Enterprise Information Systems*, 12(8-9):1034–1057, 2018.
- [69] R. Tighilt, M. Abdellatif, N. Moha, H. Mili, G. E. Boussaidi, J. Privat, and Y.-G. Guéhéneuc. On the study of microservices antipatterns: A catalog proposal. In *Proceedings of the European Conference on Pattern Languages of Programs 2020*, pages 1–13, 2020.
- [70] M. Ulrich. Addressing cascading failures. Retrieved September 2022 from <https://sre.google/sre-book/addressing-cascading-failures/>.
- [71] J. von Kistowski, S. Eismann, N. Schmitt, A. Bauer, J. Grohmann, and S. Kounev. TeaStore: A Micro-Service Reference Application for Benchmarking, Modeling and Resource Management Research. In *Proceedings of the 26th IEEE International Symposium on the Modelling, Analysis, and Simulation of Computer and Telecommunication Systems, MASCOOTS '18*, September 2018.
- [72] S. Waterworth. Stan’s robot shop – a sample microservice application. <https://www.instana.com/blog/stans-robot-shop-sample-microservice-application/>, 2018.
- [73] P. Wizenty, J. Sorgalla, F. Rademacher, and S. Sachweh. Magma: Build management-based generation of microservice infrastructures. In *Proceedings of the 11th European conference on software architecture: companion proceedings*, pages 61–65, 2017.
- [74] xiling42. Zl-ldfi with train ticket. Retrieved September 2022 from <https://github.com/xiling42/train-ticket-LDFI>, 2020.
- [75] Y. Zhang, W. Hua, Z. Zhou, G. E. Suh, and C. Delimitrou. Sinan: ML-based and qos-aware resource management for cloud microservices. In *Proceedings of the 26th ACM International Conference on Architectural Support for Programming Languages and Operating Systems*, pages 167–181, 2021.
- [76] X. Zhou, X. Peng, T. Xie, J. Sun, C. Xu, C. Ji, and W. Zhao. Poster: Benchmarking microservice systems for software engineering research. In

2018 IEEE/ACM 40th International Conference on Software Engineering:
Companion (ICSE-Companion), pages 323–324. IEEE, 2018.