

Distributed-to-Centralized: Closing the Loop on Architecture Migration via Unification Refactoring

Provakar Mondal^{*†}   

Software Innovations Lab, Virginia Tech, USA

Joshua Martin^{*}  

Software Innovations Lab, Virginia Tech, USA

Eli Tilevich[†]   

Software Innovations Lab, Virginia Tech, USA

Abstract

To address business needs and resource constraints, software systems are distributed across multiple execution sites. As requirements evolve and constraints change, distribution may need to be either increased or decreased. While increasing distribution—converting centralized software into distributed components—has been a target of a concerted research effort, decreasing distribution—unifying distributed components—has been largely neglected. Yet recent industry announcements report on reducing distribution by consolidating microservices into monoliths, albeit without a systematic treatment of the associated challenges and solutions. Hence, such migrations rely on manual, ad hoc rewrites that are difficult to plan, execute, and validate. To close the loop on architecture migration, this paper presents UNIFICATION REFACTORING, which reduces distribution by unifying distributed components into a functionally equivalent centralized system. We realize UNIFICATION REFACTORING as μ jUniter, a proof-of-concept framework that systematically transforms microservice-based architectures into modular monoliths. We evaluated μ jUniter on third-party and synthetic microservice systems. Our evaluation demonstrates that μ jUniter consistently preserves system functionality, as validated by 409 tests. It also decreases invocation latency, memory consumption, and CPU utilization by $\approx 65\%$, $\approx 72\%$, and $\approx 33\%$, respectively. Furthermore, by automating tedious and error-prone source code transformations, μ jUniter minimizes manual intervention, thereby benefiting developers seeking to reduce distribution. By closing the loop on architecture migration, UNIFICATION REFACTORING addresses a critical omission in the automatic adaptation of modern software.

2012 ACM Subject Classification Software and its engineering \rightarrow Software system structures; Computer systems organization \rightarrow Distributed architectures

Keywords and phrases Automated Refactoring, Source Code Transformation, Program Analysis, Microservices, Monoliths

Digital Object Identifier 10.4230/LIPIcs.ECOOP.2026.24

1 Introduction

Software systems continuously adapt their deployment boundaries to align with evolving operational demands, fluctuating resource constraints, and the ongoing necessity of maintenance [36]. Routine maintenance often ranges between small-scale changes to fundamental architectural transformations [47]. Some of these transformations include adapting the system’s distribution levels to cope with evolving resources and requirements. As software systems grow in size and complexity, distribution adaptation has become increasingly difficult and costly [86].

* Equal Contribution

† Corresponding Author



42 Most modern software systems follow one of the following two architectural styles: (1)
43 monoliths—larger centralized systems, or (2) microservices—distributed systems comprising
44 smaller independent services [3]. Microservices have gained substantial popularity due to
45 the growth of cloud and edge computing, thus increasing the need for larger integrated
46 systems [71]. Despite the demonstrated benefits of microservices in many modern deployments,
47 this architecture may not remain the most suitable as the system faces new requirements [43].
48 New operational realities often bring to the forefront known downsides of microservices,
49 including increased management complexity, higher latency due to network communication,
50 and additional deployment and monitoring overhead [75]. One way to address these downsides
51 is to transition the system from microservices to a monolithic architecture [25].

52 Hence, modern system architectures are being migrated from monoliths to microservices
53 and vice versa. The research community, however, has primarily focused on transforming
54 monoliths into microservices, having advanced numerous state-of-the-art program analyses
55 and transformations [32, 60, 39, 48]. In contrast, it has largely neglected the reverse journey.

56 Nevertheless, several recent industry announcements describe experiences of migrating
57 their software architectures from microservices to monoliths [41, 8, 10]. Notwithstanding
58 their fascinating experiences, these reports neither provide a systematic account of the
59 challenges involved nor describe a step-by-step process for performing the transformation.
60 Nor do they mention any automated tools that facilitated the migration. We surmise that
61 these state-of-the-practice efforts relied largely on manual rewrites that are difficult to
62 plan, execute, and verify. A systematic formulation of this migration as a refactoring—a
63 sequence of structural changes that preserve functionality—would yield new insights into this
64 architectural transformation and its challenges. Once codified, it also enables an investigation
65 into how the process could be automated.

66 This work aims to close this loop, enabling automated refactoring not only to increase
67 distribution but also to decrease it as necessary. Specifically, we present a new refactoring
68 technique—UNIFICATION REFACTORING, that transforms distributed components into a
69 semantically equivalent centralized component. We realize our approach as μ jUniter, a
70 proof-of-concept refactoring framework that takes as input **micro**(μ)-services, written using
71 the Java Spring Boot framework, and **unites** their source code into a semantically equivalent
72 monolith. To accomplish its automatic refactoring, μ jUniter introduces novel program ana-
73 lyses and code transformation routines that automate the core structural migration, thereby
74 minimizing the manual effort required to migrate microservices to monoliths. We evaluate
75 μ jUniter by applying it to third-party and synthetic microservice systems, demonstrating
76 the preservation of system functionality, reduced resource consumption, and lowered service
77 invocation latency. We also report on the automation value of μ jUniter, demonstrating that
78 it eliminates the need for manual rewrites of communication and some failure-handling logic.
79 Consequently, developer effort is shifted from low-level code refactoring to a modest oversight
80 role in policy consolidation.

81 The contributions of this paper are as follows:

- 82 1. UNIFICATION REFACTORING, an architectural refactoring for transforming the source
83 code of distributed components into a functionally equivalent centralized component.
- 84 2. μ jUniter, an automated framework that reifies UNIFICATION REFACTORING for the
85 domain of Java Spring Boot microservices.
- 86 3. An empirical evaluation of μ jUniter that demonstrates the effectiveness of the auto-
87 mated refactoring transformations and their impact on system performance and resource
88 consumption.

89 The rest of this paper is structured as follows. Section 2 provides the technical back-

90 ground and motivation for this work; Section 3 describes the system design of UNIFICATION
91 REFACTORING, its main components, and workflow; Section 4 describes how we implemented
92 μ jUniter—a proof-of-concept of our approach to automatically uniting disparate microservices
93 into a semantically equivalent monolith; Section 5 presents how we evaluated various charac-
94 teristics of μ jUniter; Section 6 discusses the applicability and limitations of our approach;
95 Section 7 describes the related state of the art; finally, Section 8 concludes the paper with
96 future work directions.

97 **2 Background and Motivation**

98 In this section, we describe the technical concepts used in this work and present real-world
99 cases that motivate this research.

100 **2.1 Background**

101 The technical background of this research includes the concepts of refactoring and the
102 architectural styles of microservices and monoliths, which we outline in turn.


103 **Refactoring**

104 In its classic definition, *refactoring* is a semantics-preserving program transformation intended
105 to improve source code quality, making it easier to maintain and extend [24]. Refactoring
106 has been applied to improve various aspects of source code and program execution [1],
107 such as increasing energy efficiency [56], bolstering security [2], adopting new language
108 features [78, 40], and migrating to new execution environments [44]. Meanwhile, most
109 contemporary software systems are distributed, with different parts of their execution running
110 on remote machines and communicating over a network. Hence, some refactoring techniques,
111 such as *partitioning* and *offloading*, transform a centralized system into a distributed execution
112 model [80, 45, 49]. However, while the literature extensively covers the decomposition of
113 software into distributed components, the inverse transformation—systematically unifying
114 distributed source code to reduce distribution—remains an unexplored frontier in automated
115 refactoring.

116 **Microservices and Monoliths**

117 The microservices architecture organizes a system into loosely coupled services that commu-
118 nicate over a network, commonly via RESTful APIs, enabling independent scaling and fault
119 isolation [81]. This architecture manages microservices via a combination of components and
120 strategies, including API gateways, load balancers, and other mechanisms. This architecture
121 promotes the independent development, deployment, and scaling of individual services,
122 making it well-suited for large-scale distributed systems [61]. However, the distributed
123 nature of microservices can incur significant operational complexity, including challenges
124 in service-to-service communication, deployment orchestration, fault tolerance, and data
125 consistency [75].

126 In contrast, a monolithic architecture consolidates functional units into corresponding
127 modules deployed as a centralized system, thereby enabling modules to interact via local
128 function calls [64]. While monoliths may offer less scaling flexibility, they can simplify
129 deployment pipelines, reduce inter-component communication latency, and lower operational
130 overhead [76].



■ **Figure 1** Migration from Microservices to Functionally Equivalent Monolithic Architecture

131 Figure 1 depicts the architectural migration from a microservices-based system, with
 132 separate microservices communicating with each other over a network, to an equivalent
 133 monolithic system, in which the original microservices become corresponding local modules,
 134 communicating with each other via local function calls.

135 2.2 Motivating Industry Announcements

136 We respond to the observation that enterprise systems sometimes need to adjust their
 137 distribution levels. Consider the *comment* to a YouTube clip on Amazon Prime Video’s
 138 migration from microservices to a monolith: “*been in IT for 30 years. it is an endless*
 139 *cycle of flipping back and forth on architecture decisions. can’t remember anymore if we*
 140 *are centralizing or decentralizing now*” [28]. As evidenced by this comment’s popularity, it
 141 expresses a shared sentiment among software developers. Our work is inspired by several
 142 prominent enterprises that, for various reasons, needed to reduce the distribution of their
 143 systems, in some cases migrating from microservices to monoliths, and subsequently released
 144 announcements reflecting on their motivations and experiences.

145 2.2.1 Prime Video

146 One announcement reports that Prime Video’s migration of a stream monitoring application
 147 reduced costs by as much as 90% [41]. As their initial serverless system grew, the team
 148 observed an unexpected spike in costs and new scaling bottlenecks. Analysis revealed that
 149 the culprits were the orchestration workflow and inter-service communication. While the
 150 announcement refers to consolidating microservices, the resulting architecture was not fully
 151 monolithic [35]. However, the migration reduced distribution and sparked many subsequent
 152 discussions about the real-world effects of overly distributed system overheads [9, 28].

153 2.2.2 Twilio Segment

154 Another announcement concerns the customer data platform—Segment. The primary function
 155 of Segment’s enterprise application is to interact with numerous external data providers. To
 156 improve fault isolation, the application was initially migrated to a microservices architecture.
 157 However, the resulting system did not meet expectations, largely due to the disadvantages of
 158 microservice complexity, prompting a reversal of the journey [8]. Specifically, the distributed
 159 microservice architecture assigned one service per provider; thus, as the number of providers
 160 increased, the number of services also increased. This growth eventually led to an inordinate
 161 increase in operational overheads and development time. Having tried various solutions to

162 address these issues, the development team determined that excessive distribution was the
163 primary cause and that the best option was to revert to the monolith.

164 2.2.3 Istio

165 The last announcement concerns Istio—a popular open-source service mesh used to manage
166 and run microservice applications [51]. This case is especially consequential: given the nature
167 of their service, the development team possessed profound expertise in microservices, yet
168 they decided to migrate their service to a monolith. The announcement first notes that
169 microservices are beneficial only when the value of independent delivery teams, rollouts, and
170 scaling exceeds the cost of orchestration. Since Istio’s control plane did not reap the benefits
171 of microservices, it was decided that the original components needed to be consolidated into
172 a monolith. This migration facilitated installation, configuration, use of VMs, maintenance,
173 scalability, and debugging. Furthermore, the development team reported that the monolith
174 reduced startup time and resource usage, resulting in improved responsiveness.

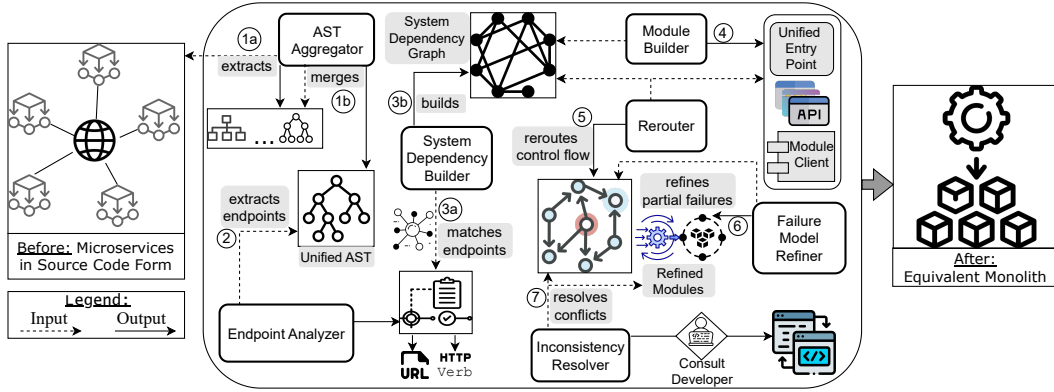
175 These announcements report on migrations performed on proprietary systems; therefore,
176 source-code-level details cannot be independently verified. Nevertheless, the described
177 business and operational realities firmly establish the necessity to reduce distribution in some
178 cases. Furthermore, none of the narratives mentions any automated tools used to facilitate
179 the migrations. Since distribution must be increased and decreased, the migration toolchain
180 must be extended accordingly. Hence, we view their experiences as motivation to develop an
181 automated refactoring that can serve as a blueprint for similar undertakings.

182 Our Perspective

183 We do not advocate for the superiority of either microservices or monolithic architectures, as
184 each possesses well-known advantages and limitations. In this paper, we describe how the
185 microservices-to-monolith migration can be framed as an automated source code refactoring,
186 and we evaluate the effectiveness of the migration and its impact on performance and resource
187 consumption. Finally, this work focuses solely on automating the source-code refactoring
188 required for the migration, assuming that transforming a microservices system into a monolith
189 has already been established as beneficial.

190 **3** System Design and Workflow

191 In this section, we describe the system design and workflow of UNIFICATION REFACTORING.
192 As an architectural refactoring, it describes a large-scale transformation that modifies
193 the software architecture while preserving the business functionality. Distributed system
194 architectures come in a wide variety, but they can benefit from UNIFICATION REFACTORING
195 when the degree of distribution needs to be reduced. Since refactoring is expected to
196 preserve the system’s software quality, UNIFICATION REFACTORING needs to maintain
197 the modularity and structural advantages of distributed components while eliminating the
198 inherent communication and orchestration overheads. Hence, it primarily focuses on replacing
199 remote inter-component communication with local function calls. We demonstrate the main
200 components of UNIFICATION REFACTORING by transforming a microservice system (Before)
201 into a modular monolith (After), where each original service is preserved as a distinct module,
202 connected via local calls, with a unified entry point. Figure 2 depicts these components and
203 the general flow between them, which we describe next.



■ Figure 2 UNIFICATION REFACTORING: Main Components and General Flow

204 3.1 Abstract Syntax Tree (AST) Aggregator

205 An Abstract Syntax Tree (AST) is a common representation of a program’s source code used
 206 for various language processing tasks, such as static analysis and program transformation [59].
 207 In this structured tree-based representation, each node corresponds to a syntactic construct,
 208 such as a class, method, or statement. AST preserves the code’s syntactic and structural
 209 information, while abstracting low-level details. Having received the microservice source
 210 code, the *Aggregator* first extracts the AST for each microservice, (1a). The *AST Aggregator*
 211 then merges the ASTs of all microservices into a single unified AST, (1b). The aggregation
 212 strategy maintains service boundaries by placing each service into a separate module. In
 213 cases where duplicate names arise among the aggregated microservices, naming clashes are
 214 resolved by renaming the duplicates. Transforming at the AST level enables reasoning about
 215 and manipulating code in a language-aware way, ensuring that transformations preserve
 216 semantics and produce syntactically correct output.

217 Formally, let the set of microservices be $\mu\mathcal{S} = \{\mu S_1, \mu S_2, \dots, \mu S_n\}$ and their ASTs be
 218 $\mathcal{A} = \{AST_1, AST_2, \dots, AST_n\}$ with $AST_i \leftarrow \text{Extract}(\mu S_i)$. The aggregation is defined as:
 219 $AST_{\mathcal{U}} = \bigoplus_{i=1}^n AST_i$.

220 3.2 Endpoint Analyzer

221 Once the unified AST ($AST_{\mathcal{U}}$) is in place, the *Endpoint Analyzer* traverses each module to
 222 identify its server endpoints (interfaces the module provides) and client endpoints (interfaces
 223 it interacts with). Endpoints include attributes such as URL, HTTP method (POST, PUT,
 224 etc.), and payload (JSON, XML, or plain text). The *Endpoint Analyzer* records the location
 225 of each endpoint in the AST for later transformation, (2).

226 For each module M_i in $AST_{\mathcal{U}}$, identify:

$$227 E_i^{\text{server}} = \{(\text{url}(e_s), \text{verb}(e_s), \text{payload}(e_s)) \mid e_s \in AST(M_i), e_s \text{ is a server endpoint}\}$$

$$228 E_i^{\text{client}} = \{(\text{url}(e_c), \text{verb}(e_c), \text{payload}(e_c)) \mid e_c \in AST(M_i), e_c \text{ is a client endpoint}\}$$

230 3.3 System Dependency Builder

231 This component builds a System Dependency Graph (SDG) from the endpoint information.
 232 An SDG is a graph in which nodes represent modules (former microservices), and edges
 233 represent calls from client endpoints to their corresponding server endpoints. By matching

234 endpoints based on their attributes (URL, HTTP verb, and payload), (3a), this *Builder*
 235 constructs the SDG that captures the complete set of inter-service communication pathways,
 236 (3b). This graph serves as a blueprint for the subsequent transformation, ensuring that
 237 all relevant communication will be rewired into local calls. SDG also provides a global
 238 architectural view that can be visualized for examination and debugging before refactoring.

239 Define the internal communication set:

240 $\text{IntComm} = \{(e_c, e_s) \mid e_c \in E_i^{\text{client}}, e_s \in E_j^{\text{server}}, i \neq j, \text{Match}(e_c, e_s)\}$, where

241 $\text{Match}(e_c, e_s) \Rightarrow \text{url}(e_c) = \text{url}(e_s) \wedge \text{verb}(e_c) = \text{verb}(e_s) \wedge \text{payload}(e_c) = \text{payload}(e_s)$.

243 The condition $i \neq j$ excludes intra-module calls, which are already local. UNIFICATION
 244 REFACTORING specifically targets transforming distributed (remote) calls into centralized
 245 (local) ones. The System Dependency Graph is then defined as: $\text{SDG} = (V, E)$, where

246 $V = \{M_1, \dots, M_n\} \wedge E = \{(M_i, M_j) \mid \exists(e_c, e_s) \in \text{IntComm}, e_c \in M_i, e_s \in M_j\}$.

247 3.4 Module Builder

248 The SDG in the previous step contains all the modules of the future monolith. However,
 249 these modules must be implemented in code, which is the purpose of the *Module Builder*
 250 component. It begins by establishing a centralized entry point for the monolith. In a
 251 distributed microservice system, each service maintains its own entry point and initialization
 252 logic, along with configuration for functionality such as security and integration with external
 253 systems. To unify these microservices, this *Builder* consolidates their entry points into a
 254 single main entry point that represents the entire system. This unified entry point preserves
 255 all essential configurations and behaviors from each constituent service. It also ensures that
 256 they coexist without conflict within a unified monolith. As a result, the resulting system can
 257 act as a single cohesive unit that preserves the semantics of the original microservices.

258 With the main entry point complete, the *Builder* creates a Module API for each module,
 259 thereby exposing its functionality to other modules. This step also bridges the control flow
 260 between modules, beginning with the server endpoint's location within the AST. These APIs
 261 expose a proxy that handles incoming requests and prepares responses, enabling modules to
 262 communicate with one another by calling these APIs.

263 Once the APIs are in place, the *Builder* produces Module Clients for the modules that
 264 invoke these APIs. Each client encapsulates the logic required to invoke another module's
 265 API, taking over the role that remote communication played in the microservice architecture.
 266 In addition to calling the API methods, module clients handle data adaptation to maintain
 267 compatibility when different modules use distinct but structurally similar data representations.
 268 In a distributed system, such adaptation often occurs implicitly during serialization and
 269 deserialization [79]; in a monolith, it must be handled explicitly to ensure that data exchanged
 270 between modules remains correct, consistent, and isolated from unintended side effects [53].
 271 This explicit adaptation is also required to prevent unintended coupling between modules'
 272 internal states, so that they remain logically separated, as in the original distributed system.

273 In summary, the *Module Builder* constructs the unified entry point, module API, and
 274 module clients (4) for the resulting monolith. Formally, all microservice-specific entry points
 275 $\{EP_1, EP_2, \dots, EP_n\}$ are consolidated into a single unified entry point $EP_{\mathcal{U}}$, preserving their
 276 initialization logic and configurations:

277 $EP_{\mathcal{U}} = \text{Unify}(EP_1, EP_2, \dots, EP_n)$.

278 Then, for each server endpoint $e_s \in E_i^{\text{server}}$, the *Builder* creates an API method $API_i(e_s)$,
 279 which exposes the corresponding functionality as an internal callable:

$$280 \quad API_i = \{f_{e_s} \mid e_s \in E_i^{\text{server}}\}, \quad f_{e_s} : \text{payload}(e_s) \mapsto \text{response}(e_s).$$

281 Finally, for each client endpoint $e_c \in E_i^{\text{client}}$ that matches a server endpoint $e_s \in E_j^{\text{server}}$
 282 (with $i \neq j$), the *Builder* generates a client method $Client_i(e_c)$ that invokes the corresponding
 283 API method f_{e_s} while applying an adaptation function α if the payload or response structures
 284 differ:

$$285 \quad Client_i(e_c) = \lambda x. f_{e_s}(\alpha(x)), \text{ where } \alpha : \text{payload}(e_c) \rightarrow \text{payload}(e_s).$$

286 3.5 Rerouter

287 The *Rerouter* component replaces previously invoked remote endpoints of microservices with
 288 calls to the corresponding local module methods, ⑤. The original microservices used remote
 289 communication across the network. To reroute control flow from this remote communication
 290 over the network to local function calls, the *Rerouter* replaces all call sites that initiate
 291 the remote communication with the newly built module clients. The rerouting transforms
 292 distributed control flow into local, intra-process control flow.

293 Formally, after rerouting, the client of module M_i invokes the API of module M_j that
 294 corresponds to the server endpoint e_s of the original microservice system:

$$295 \quad \forall (e_c, e_s) \in \text{IntComm} : \underbrace{\text{Call}(e_c, x)}_{\text{remote}} \xrightarrow{\text{Reroute}} \underbrace{\text{Call}(f_{e_s}, \alpha(x))}_{\text{local, intra-process}}, \quad f_{e_s} = API_j(e_s),$$

296 where $x \in \text{payload}(e_c)$ denotes the argument (payload data) passed at the original client call
 297 site, and α is the adaptation function that maps the payload expected by the client endpoint
 298 e_c into the payload schema required by the server endpoint e_s .

299 3.6 Failure Model Refiner

300 To preserve semantic equivalence, UNIFICATION REFACTORING must ensure that the
 301 architectural migration accounts for the disparate failure models of distributed and centralized
 302 executions. Distributed systems are prone to *partial failures*, conditions under which a sub-
 303 component becomes unreachable or inoperational while the rest of the system continues
 304 to function normally [55]. In contrast, the failure model of monoliths is *total* rather than
 305 *partial*. In other words, a monolithic system can also experience faults in its business logic
 306 and other local components, and if these are not recoverable, the system terminates [52].
 307 The failures that pertain to distributed execution can no longer occur once the distributed
 308 components have been consolidated into a monolith, so the logic for handling such failures
 309 becomes superfluous. To remove this now-unnecessary logic, UNIFICATION REFACTORING
 310 employs a *Failure Model Refiner*, whose high-level functionality is described next, ⑥.

311 Formally, let e_s be a server endpoint and $x \in \text{payload}(e_s)$. In a distributed environment,
 312 the result of an invocation by e_c , denoted \mathcal{O}_{dist} :

$$313 \quad \mathcal{O}_{dist}(e_c, x) \in \{\text{response}(e_s)\} \cup \mathcal{E}_{app} \cup \mathcal{E}_{dist}$$

314 where \mathcal{E}_{app} represents application-level exceptions (e.g., business logic errors) and \mathcal{E}_{dist}
 315 represents distribution-induced failures (e.g., timeouts, network partitions).

316 By unifying distributed components, UNIFICATION REFACTORING eliminates the network
 317 as a source of failure. The refactoring process maps the original exception handlers H

318 via a transformation function $\mathcal{T}(H)$. If H catches an element of \mathcal{E}_{dist} , it is flagged as
 319 *distribution-induced dead code*, which should either be pruned entirely or perhaps substituted
 320 with local fallback logic. Consequently, the monolithic outcome \mathcal{O}_{mono} is a refinement where
 321 $\mathcal{O}_{mono} = \mathcal{O}_{dist} \setminus \mathcal{E}_{dist}$. Hence, the role of the *Failure Model Refiner* is to remove only
 322 the failure-handling logic that is no longer possible due to the replacement of distributed
 323 communication with local functions. The original business logic, which should remain intact,
 324 can experience other failures, whose handling logic must be preserved.

325 3.7 Inconsistency Resolver

326 Finally, the *Inconsistency Resolver* component addresses cross-cutting consolidation issues
 327 that arise from unifying multiple microservices into a monolith, (7). The resolution involves
 328 resolving naming conflicts, internalizing API gateway routes to prevent endpoint collisions,
 329 and merging security policies so that the monolith enforces the strictest applicable rules.
 330 This step ensures that the monolith is both functionally correct and operationally coherent.
 331 Unfortunately, the resolution does not easily lend itself to complete automation, as many
 332 of these non-functional concerns are quite domain-specific. For example, security policies
 333 are known to be defined and enforced in ways that vary widely across applications and even
 334 deployments. Hence, both manual examination and correction are required.

335 At the end of this process, the refactoring produces a modular monolith whose external
 336 behavior is indistinguishable from that of the original microservice system, but which
 337 executes without inter-service network calls. The transformed system continues to adhere to
 338 well-established software engineering principles, including high cohesion, low coupling, and
 339 statelessness. The refactoring transformations typically reduce aggregate invocation latency
 340 by minimizing communication overhead, thereby making remote inter-service communication
 341 equivalent to local inter-module calls. Additionally, orchestration overhead decreases due
 342 to the reduced need for service discovery, gateways, internal traffic management, and other
 343 management services.

344 Formally, the resulting monolith is defined as:

$$345 \mathcal{M} = (\{M_1, \dots, M_n\}, EP_{\mathcal{U}}),$$

346 where each module $M_i = (API_i, \{Client_i(e_c) \mid e_c \in E_i^{client}\})$ corresponds to a former
 347 microservice μS_i , and all inter-service communication is transformed from distributed network
 348 calls into intra-process function calls. Interaction between modules is then expressed as:

$$349 \forall e_c \in E_i^{client}, e_c \mapsto e_s \in E_j^{server}, x \in \text{payload}(e_c) : \\ 350 M_i.Client_i(e_c)(x) \rightarrow M_j.API_j(e_s)(x) \in \mathcal{O}_{mono}(e_c, x)$$

351 where $\mathcal{O}_{mono}(e_c, x) = \{\text{response}(e_s)\} \cup \mathcal{E}_{app}$. By explicitly excluding \mathcal{E}_{dist} (fault-handling
 352 logic required for distribution-induced failures), UNIFICATION REFACTORING fully transforms
 353 a distributed inter-component communication into semantically equivalent intra-process
 354 execution.

355 4 Implementation: μ JUniter

356 As a concrete implementation of UNIFICATION REFACTORING, we created μ JUniter that
 357 targets microservice applications built with the Java Spring Boot framework [7]. Both
 358 empirical and practical reasons motivate this choice. On the empirical side, microservices are
 359 mainstream in industry: 37% of programmers report working with microservices, and 34% of

360 them use Java to do so, according to the 2022 JetBrains Developer Ecosystem survey [37].
361 Within the Java ecosystem, Spring Boot dominates microservice development, with 59%
362 of Java developers reporting that they use this framework, according to the 2023 JRebel
363 Java Productivity Report [67]. These statistics highlight that Java and Spring Boot are not
364 only widely used but also representative platforms for enterprise-grade microservices. On the
365 practical side, Spring Boot’s heavy reliance on annotations and declarative configurations
366 provides a non-trivial and illustrative testbed [57]. Hence, we decided to focus on Spring
367 Boot microservices as the target of our implementation. By demonstrating the ability to
368 unify these microservices, `µjUniter`’s design reveals implementation insights that can be
369 adapted to other languages and frameworks. Recall that the system design of UNIFICATION
370 REFACTORING comprises seven main components, whose implementation in `µjUniter` is
371 described next.

372 4.1 AST Aggregator

373 For all source code analysis and transformation tasks, we used the Spoon Java library [66].
374 This library parses each microservice’s source code into a Spoon’s `CtModel` AST, which
375 provides an API for traversing, editing, and regenerating source code. For each extracted
376 microservice AST, the *Aggregator* resolves naming conflicts at the package and class level
377 and aggregates them into a single unified tree.

378 This aggregation requires certain nuance, as a naïve aggregation of all microservice root
379 packages into one may create conflicting fully qualified names. Such conflicts may cause
380 undefined behavior later, when the AST is converted back into source code. Hence, the
381 *Aggregator* must ensure that the fully qualified names of all classes are unique. Possible
382 solutions include class renaming [72], modular packaging[42], or various modularization
383 approaches such as Java 9 modules [14], Maven Modules [23], or Spring Modulith [17].
384 While the modularization approaches provide the best encapsulation, they introduce new
385 dependencies, increase complexity, and are inconsistent across programming languages [5].
386 For this reason, we use modular packaging, placing each service’s source code in a package
387 named after the service. This step first places each service into a uniquely named package,
388 then aggregates all packages into a single AST.

389 This co-location of all code in a single representation is essential for whole-system
390 analysis, enabling the detection of cross-service relationships and preventing high coupling
391 in the services’ internal logic. Algorithm 1 shows the step-by-step pseudo-code of the AST
392 aggregation task.

393 4.2 Endpoint Analyzer

394 For this component, we implement AST traversal routines that locate service endpoints
395 by detecting Spring-specific Java annotations. Server endpoints are identified through
396 annotations such as `@RestController` and `@RequestMapping` (including derived forms like
397 `@GetMapping` and `@PostMapping`), while client endpoints are identified via annotations such
398 as `@FeignClient`. These annotations and their programming conventions serve as reliable
399 markers for identifying endpoints. For each endpoint, we record not only the metadata (URL,
400 HTTP verb, payload) but also the exact AST node location, ensuring that transformations in
401 later stages can directly modify or replace the corresponding methods. This annotation-driven
402 analysis shows how tightly implementation details are coupled to framework conventions,
403 and reinforces that support for other ecosystems would require re-implementing this analysis
404 step with their respective idioms.

Algorithm 1 AST Extraction and Aggregation

Input: $\mu S = \{\mu S_1, \dots, \mu S_n\}$: Set of Input Microservices
Output: AST_U : Unified Abstract Syntax Tree

```

1  $\mathcal{A} \leftarrow \emptyset$ 
2 forall  $\mu S_i \in \mu S$  do
3    $AST_i \leftarrow \text{Spoon.ExtractAST}(\mu S_i)$ 
4    $AST_i \leftarrow \text{ApplyModularPackaging}(AST_i, \text{packageName} = \mu S_i.\text{name})$ 
   /* Compare with previously aggregated ASTs to ensure unique class names */
5    $AST_i \leftarrow \text{ResolveNamingConflicts}(AST_i, \mathcal{A})$ 
6    $\mathcal{A} \leftarrow \mathcal{A} \cup \{AST_i\}$ 
7  $AST_U \leftarrow \text{AggregateASTs}(\mathcal{A})$ 
8 return  $AST_U$ 

```

405 **4.3 System Dependency Builder**

406 This component constructs a System Dependency Graph (SDG) using Algorithm 2 by match-
 407 ing endpoints in the unified AST. The built SDGs are stored in Neo4j [46], a graph database,
 408 via the `NeoRepository` class. Neo4j enables not only persistence but also interactive
 409 exploration through tools such as Neo4j Bloom [34]. Before proceeding with a refactoring
 410 transformation, a developer may want to validate that all inter-service communication paths
 411 are captured. Although visualizations are not strictly necessary to implement `µJUniter`,
 412 modern developers expect refactoring tools to facilitate understanding and validation. As
 413 an example, Figure 3 visualizes the built SDGs for our evaluation subjects prior to their
 414 unification. We detail the subjects in Section 5.1. In these diagrams, the red and blue circles
 415 represent the microservices and their endpoints, respectively.


Algorithm 2 Endpoint Extraction, Matching, and SDG Construction

Input: AST_U : Unified AST
Output: $SDG = (V, E)$: System Dependency Graph
Output: `IntComm`: Internal Communication Set

```

1  $V \leftarrow \{M_1, \dots, M_n\} \leftarrow \text{ExtractModules}(AST_U)$ 
2  $\text{IntComm} \leftarrow \emptyset$ 
3 foreach  $M_i \in V$  do
4    $E_i^{server} \leftarrow \text{LocateEndpoints}(M_i, \{\text{@RestController}, \text{@RequestMapping}, \dots\})$ 
5    $E_i^{client} \leftarrow \text{LocateEndpoints}(M_i, \{\text{@FeignClient}\})$ 
6 foreach  $e_c \in E_i^{client}$ ,  $e_s \in E_j^{server}$ ,  $i \neq j$  do
7   if  $\text{url}(e_c) = \text{url}(e_s) \wedge \text{verb}(e_c) = \text{verb}(e_s) \wedge \text{payload}(e_c) = \text{payload}(e_s)$  then
8      $\text{IntComm} \leftarrow \text{IntComm} \cup \{(e_c, e_s)\}$ 
9  $E \leftarrow \{(M_i, M_j) \mid \exists (e_c, e_s) \in \text{IntComm}, e_c \in M_i, e_s \in M_j\}$ 
10  $SDG \leftarrow (V, E)$ 
11 return  $(SDG, \text{IntComm})$ 

```



■ **Figure 3** System Dependency Graph Exemplars

416 4.4 Module Builder

417 This component performs three actions: (1) creates a unified main entry point, (2) creates
 418 Module APIs, and (3) builds Module Clients. (1) A new main class is created to consolidate the
 419 entry points of all services. This unified class merges `@SpringBootApplication` annotations,
 420 initialization logic, and configuration beans, while resolving conflicts through systematic
 421 renaming and refactoring. In Java, this requires preserving all in-code configurations from
 422 each service’s main class, including annotations, beans, and methods for enabling clients,
 423 defining APIs, implementing security, and configuring repositories. The result is a single
 424 monolithic entry point that preserves the startup semantics of all original microservices. (2)
 425 Module APIs are generated for accessing module functionality. These APIs act as proxies
 426 for controller methods, creating explicit entry points for inter-module calls and defining a
 427 controlled location for inter-module communication. (3) Module Clients are generated for
 428 invoking module APIs. Algorithm 3 details these three actions.

■ **Algorithm 3** Module Building for Unified Entry Point, APIs, and Clients

Input: $AST_{\mathcal{U}}, SDG, \text{IntComm}$
Output: $\mathcal{M} = \{M_1, \dots, M_n\}$: Set of Monolith Modules

```

/* unifying microservice-specific entry points (EPs) into one */
1  $EP_{\mathcal{U}} \leftarrow \text{Unify}(EP_1, \dots, EP_n)$ 
2  $\mathcal{M} \leftarrow \emptyset$ 
3 foreach  $M_i \in SDG.V$  do
4    $API_i \leftarrow \{f_{e_s} \mid e_s \in E_i^{server}\}$ , with  $f_{e_s} : \text{payload}(e_s) \mapsto \text{response}(e_s)$ 
5    $Client_i \leftarrow \emptyset$ 
6   foreach  $(e_c, e_s) \in \text{IntComm}$ ,  $e_c \in M_i$  do
7     /*  $\alpha : \text{payload}(e_c) \rightarrow \text{payload}(e_s)$  is the adaptation function */
8      $Client_i(e_c) \leftarrow \lambda x. f_{e_s}(\alpha(x))$ 
9    $\mathcal{M} \leftarrow \mathcal{M} \cup \{M_i(API_i, Client_i)\}$ 
9 return  $\mathcal{M}$ 

```

429 Each client method invokes the corresponding API method in the target module and also
 430 provides an additional routine, deep-copy, for type conversion [85]. This routine emulates
 431 structural typing within a nominally typed language for the microservice parameters and
 432 return types that are Java type-incompatible. Accommodating such type-incompatibility
 433 requires generating adaptation logic. When transferring data across distributed sites, struc-
 434 turally compatible but type-unrelated objects (parameters and return types) can be serialized
 435 into a single type and then unserialized into another [68]. When the distribution is removed,
 436 this component emulates serialization via deep copying. Specifically, deep-copy routines are

437 generated for all pairs of incompatible types whose objects need to be assigned to each other.

438 The deep copy is generated as follows. Immutable objects are passed directly, iterables are
 439 reconstructed item by item, and custom models are recreated via constructors and accessors,
 440 with fields copied using getter and setter methods. Because server endpoints can serve
 441 multiple clients with different expectation types, the copy-conversion logic is generated in
 442 the client rather than the API class. If automated copying fails, serialization is retained as a
 443 conservative fallback.

444 4.5 Rerouter

445 This component replaces the remote calls in client modules with local calls to the generated
 446 Module Client methods. The replacement is performed via AST rewriting: each method
 447 invocation node is replaced with a new node corresponding to the local client call, preserving
 448 its arguments and signature. Spoon's ability to query the AST for invocations by annotation
 449 and type makes this possible without fragile text-based pattern matching, such as scanning
 450 source files for occurrences of `RestTemplate.getForObject(...)` or `FeignClient` method
 451 calls, or searching for string literals like `"http://<service>/<endpoint>"`. This rerouting
 452 step completes the replacement of network calls with local calls. Algorithm 4 shows the
 453 step-by-step pseudo-code of the control flow rerouting task.

■ Algorithm 4 Control Flow Rerouting

Input: \mathcal{M} , IntComm
Output: $\mathcal{M}_{rerouted}$: Modules with Local Method Invocations

```

/* Initialize the rerouted set with the original unified modules */
1  $\mathcal{M}_{rerouted} \leftarrow \mathcal{M}$ 
2 foreach  $(e_c, e_s) \in \text{IntComm}$  do
3    $M_i \leftarrow$  module owning  $e_c \in \mathcal{M}_{rerouted}$ 
4    $M_j \leftarrow$  module owning  $e_s \in \mathcal{M}_{rerouted}$ 
   /* Rewrite AST: replace Feign/RestTemplate calls with direct local calls */
5    $M'_i \leftarrow \text{ReplaceInvocationAST}(e_c, \text{call} = M_j.API_j(e_s), \text{args} = \alpha(\text{payload}(e_c)))$ 
6    $\mathcal{M}_{rerouted} \leftarrow (\mathcal{M}_{rerouted} \setminus \{M_i\}) \cup \{M'_i\}$ 
7 return  $\mathcal{M}_{rerouted}$ 

```

454 4.6 Failure Model Refiner

455 In the Java Spring Boot ecosystem, partial failures are typically managed through three
 456 primary mechanisms: (1) low-level network exceptions from the `java.net` hierarchy, (2)
 457 declarative resilience patterns such as Spring Cloud Circuit Breaker or `Resilience4j`, and
 458 (3) HTTP status code mapping via `FeignException` or `HttpStatusException`. To im-
 459 plement the transformation function $\mathcal{T}(H)$ introduced in Section 3.6, `µJUniter` systematically
 460 identifies these artifacts within the Spoon AST and determines if they represent business
 461 logic or distribution remnants.

462 `µJUniter` identifies distribution-induced failures (\mathcal{E}_{dist}) by scanning the unified AST for spe-
 463 cific types and annotations that lose their semantic usefulness in a local context. These include
 464 connection failures (e.g., `ConnectException`, `SocketTimeoutException`), framework-specific
 465 RPC exceptions (e.g., `ResourceAccessException`, `RetryableException`), and resilience an-
 466 notations (e.g., `@CircuitBreaker`, `@Retry`). Once identified, `µJUniter` applies domain-specific

■ **Algorithm 5** Failure Model Refinement and Handler Transformation

```

Input:  $\mathcal{M}_{rerouted}$ 
Output:  $\mathcal{M}_{refined}$ : Modules with transformed failure handling
1  $\mathcal{E}_{dist} \leftarrow \{\text{ConnectException}, \text{FeignException}, \text{TimeoutException}, \dots\}$ 
2  $\mathcal{A}_{dist} \leftarrow \{\text{@CircuitBreaker}, \text{@Retry}, \text{@Bulkhead}\}$ 
3  $\mathcal{M}_{refined} \leftarrow \mathcal{M}_{rerouted}$  // Initialize with rerouted modules
4 foreach  $M_i \in \mathcal{M}_{refined}$  do
    /* Identify and refine resilience annotations ( $\mathcal{A}_{dist}$ ) */
5      $nodes \leftarrow \text{Spoon.findAnnotated}(M_i, \mathcal{A}_{dist})$ 
6     foreach  $n \in nodes$  do
7         if  $n.hasFallback()$  then
8             |  $\text{ApplyBusinessFallback}(n)$  // Preserve functional defaults
9             |  $\text{Spoon.removeAnnotation}(n)$ 
        /* Transform exception handlers  $H$  via function  $\mathcal{T}(H)$  */
10     $handlers \leftarrow \text{Spoon.findTryCatchBlocks}(M_i)$ 
11    foreach  $H \in handlers$  do
12        if  $H.getCaughtType() \in \mathcal{E}_{dist}$  then
13            |  $sideEffects \leftarrow \text{AnalyzeSideEffects}(H)$ 
14            | if  $sideEffects \in \{\text{Logging}, \text{Retry}\}$  then
15                |  $\mathcal{T}(H) \leftarrow \emptyset$  // Prune distribution-induced dead code
16            | else
17                |  $\mathcal{T}(H) \leftarrow \text{RefactorToLocal}(H)$  // Map to application logic
        /* Update  $M_i$  and remove obsolete checked exception signatures */
18     $M'_i \leftarrow \text{CleanupCheckedExceptions}(M_i, \text{"java.io.IOException"})$ 
19     $\mathcal{M}_{refined} \leftarrow (\mathcal{M}_{refined} \setminus \{M_i\}) \cup \{M'_i\}$ 
20 return  $\mathcal{M}_{refined}$ 

```

467 transformation rules based on the functionality of the original handlers. A significant chal-
468 lenge arises when a developer implements advanced business-logic strategies, such as handling
469 network failures with a local fallback, e.g., returning a default cached object when a service
470 is unreachable. If $\mu\text{JUniter}$ finds a `catch` block for an element of \mathcal{E}_{dist} containing such logic,
471 it preserves the business response but refactors the trigger to a local equivalent. It is worth
472 noting that $\mu\text{JUniter}$ does not always correctly disentangle these distributed and business
473 functionalities, a lack of precision that we consider acceptable for a refactoring tool typically
474 used under the programmer's oversight.

475 For resilience patterns, $\mu\text{JUniter}$ removes the proxying infrastructure because it imposes
476 unnecessary overhead in the absence of distributed faults. Specifically, maintaining a
477 half-open [77] state for a local method call would introduce unnecessary latency without
478 providing fault isolation. Hence, such patterns need to be removed to ensure efficient local
479 execution. A final practical hurdle is the handling of checked exceptions. Many microservices
480 declare `throws IOException` solely to account for network volatility, a condition that can
481 no longer be present once distributed communication is removed. $\mu\text{JUniter}$ addresses this
482 now-unnecessary exception declaration by performing a *signature cleanup*. It traverses the
483 call graph in the modular monolith and removes distribution-related checked exceptions from
484 method signatures if the new local implementation no longer throws them. This cleanup

485 ensures that the resulting monolith appears to have been originally engineered for centralized
 486 execution. Algorithm 5 details the implementation of this refinement process.

487 4.7 Inconsistency Resolver

488 As the final step, this component addresses cross-cutting consolidation issues. For dependency
 489 management, we rely on Maven [22], which provides both a build system and a consistent
 490 mechanism for reconfiguring the project as a single modular monolith. To address name
 491 conflicts in dependency-injection beans, we extend our Spoon-based transformations to
 492 rename beans and update their references. For security consolidation, we analyze the Spring
 493 Security configurations in each service and merge them into the monolith, defaulting to the
 494 strictest applicable policy. Path rewrites introduced by API gateways necessitate endpoint-
 495 path normalization to avoid collisions. These tasks highlight that, although architectural
 496 refactoring is largely automated, system-specific configurations still require careful handling
 497 and, in some cases, human oversight.

498 This component’s functionality can be described as a sequential pipeline of refinements
 499 that resolves cross-cutting discrepancies across four distinct tasks:

500 **Task I: Project Reconfiguration** ($\mathcal{M}_{refined} \rightarrow \mathcal{M}_{struct}$) By leveraging Maven, this re-
 501 finement establishes a unified build lifecycle. It consolidates independent build scripts
 502 into a multi-module project structure, reconciling top-level dependency management to
 503 ensure a consistent classpath across the unified environment.

504 **Task II: Dependency Injection Namespace Resolution** ($\mathcal{M}_{struct} \rightarrow \mathcal{M}_{beans}$) To prevent
 505 bean name collisions in the shared Spring context, we apply Spoon-based transforma-
 506 tions to programmatically rename conflicting beans. All associated injection points and
 507 references are updated across the codebase to ensure runtime stability.

508 **Task III: Security Policy Synthesis** ($\mathcal{M}_{beans} \rightarrow \mathcal{M}_{sec}$) For each constituent service, the
 509 configurations of `WebSecurityConfigurerAdapter` are analyzed to produce a unified
 510 security context, defaulting to the strictest applicable CORS/CSRF policies while surfacing
 511 complex semantic overlaps for developer oversight.

512 **Task IV: Endpoint Path Normalization** ($\mathcal{M}_{sec} \rightarrow \mathcal{M}_{monolith}$) Finally, API gateway rout-
 513 ing logic is reconciled by normalizing endpoint paths. Controller mappings are rewritten
 514 to include the root paths previously managed by the gateway, preventing URL collisions
 515 and ensuring internal communication parity with the original distributed system.

516 5 Evaluation

517 The following research questions drive the evaluation of `µjUniter`:

- 518 1. **RQ1: Effectiveness:** How effectively does `µjUniter` reify UNIFICATION REFACTORING
 519 in transforming microservices into an equivalent monolith?
- 520 2. **RQ2: Performance:** How `µjUniter`’s migration affect application performance charac-
 521 teristics?
- 522 3. **RQ3: Automation Value:** What is the value of `µjUniter` automating the required
 523 source code transformations?

524 We begin by describing our evaluation subjects and then present our experimental results,
 525 which address the aforementioned questions.

5.1 Evaluation Subjects

Recall that μ jUniter targets Spring framework microservice applications. For the evaluation, we prioritized contemporary Spring Boot applications from different domains, including a booking system, a financial advisor, and a cloud platform, that can serve as representative benchmarks. Specifically, we selected the following three third-party Spring applications. Train Ticket [27] and Piggy Metrics [50]—applications whose backends are structured as microservices. The third example, which we refer to as SCShowcase (an abbreviation of SpringCloudShowcase), is a recent microservice tutorial that demonstrates the advantages of Spring Cloud as a microservice platform [54]. To stress-test the applicability of μ jUniter to microservice systems containing large numbers of microservices and their interactions, we created a generator that automatically synthesizes such applications. The generator is parameterized by the number of microservices and their interaction probabilities. We describe our evaluation subjects in greater detail next.

Train Ticket

Train Ticket is a dedicated microservice benchmark that implements a complete train-ticket booking system with a web-based front end. This benchmark has been widely used in prior research [82, 11, 74, 12]. The Train Ticket comprises 41 microservices: 37 Java Spring Boot, 2 Python, 1 JavaScript, and 1 Go (incomplete in the repository). Because μ jUniter only supports Java Spring microservices, we could not apply it to the other three services implemented in Python and JavaScript. Although we could have excluded them from evaluation, we instead chose to manually translate them to Java because of their relative simplicity and reliance on external libraries with Java equivalents. This system serves as a realistic substitute for an enterprise system, given the complexity of service interactions and the large number and diversity of services.

Piggy Metrics

Piggy Metrics is a financial advisor application. While primarily intended as a tutorial, this system has been used as a benchmark for testing and monitoring microservices [20]. The business logic is contained in the accounts, notifications, and statistics services. The only change made to Piggy Metrics before applying μ jUniter was the removal of references to an obsolete external API, causing runtime errors.

SCShowcase

This subject's microservices are implemented in Java 21 using the latest Springboot version. The business logic appears in the employee, department, and organization services. Each service manages its own domain data and exposes REST APIs, while service discovery, an API gateway, and configuration management provide the supporting infrastructure.

RM*

To create complex microservice systems, we developed a Random Microservice Generator, parameterized by the number of microservices and their interaction probabilities. Each generated microservice includes a single controller endpoint and has a certain probability of calling other microservices within the system. Because microservice systems typically

566 exclude cyclic dependencies [18], the generated microservices must avoid interactions that
 567 would create a cycle. To implement this policy, our generator uses a directed acyclic graph.

568 For our evaluation, we synthesized and unified three microservice systems. RM-10-0.2 is a
 569 synthetic microservice system with 10 services and a 0.2 service interaction rate; RM-10-1.0
 570 increases the service interaction rate; and RM-100-0.02 increases the number of services.
 571 These properties are specifically designed to stress-test μ JUniter with respect to the number
 572 of services and endpoints.

573 Table 1 provides an overview of these evaluation subjects. For some subjects, the number
 574 of client and server endpoints differs. This discrepancy arises because some leaf-node services
 575 are pure callers with no exposed server endpoints.

■ **Table 1** Evaluation Subjects: Microservice System Details

| Subject Name | No. of Microservices | No. of Server Endpoints | No. of Client Endpoints | Endpoint Library |
|---------------|----------------------|-------------------------|-------------------------|-------------------|
| Train Ticket | 40 | 98 | 157 | Spring RESTClient |
| Piggy Metrics | 3 | 4 | 4 | Spring OpenFeign |
| SCShowcase | 3 | 13 | 4 | Spring OpenFeign |
| RM-10-0.2 | 10 | 8 | 14 | Spring WebClient |
| RM-10-1.0 | 10 | 9 | 45 | Spring WebClient |
| RM-100-0.02 | 100 | 84 | 179 | Spring WebClient |

576 5.2 RQ1: Effectiveness

577 For a refactoring to be effective, it must modify the program’s structure as intended while
 578 preserving the original application’s functionality. While we verified that the microservices
 579 were transformed into a monolith through a straightforward source-code examination, we
 580 sought systematic validation that the functionality was preserved. Although refactoring
 581 was originally introduced as “a semantics-preserving transformation,” [63], verifying that a
 582 semantics is indeed preserved would require first defining it formally. In the absence of formal
 583 methods that scale to microservices, we resort to testing as a pragmatic, evidence-based
 584 approach to demonstrate functional preservation. In other words, in this evaluation, we
 585 verify if the *before* (i.e., original microservices) and *after* (i.e., refactored monolith) pass
 586 the same test cases. Of course, Dijkstra’s dictum that “testing does not prove the absence
 587 of bugs, only their presence” [16] still stands, but testing remains the most widely applied
 588 option for verifying the preservation of functionality in practical settings [58, 38].

589 For our evaluation, we first applied μ JUniter to all our evaluation subjects, producing a
 590 modular monolith version for each. We then tested both the *before* and *after* versions and
 591 verified that the test suites produced the same results. Whenever possible, we used test
 592 suites that either came with the evaluation subjects or were adapted from prior research [74].
 593 In the absence of tests, we conducted our own testing of the main functionalities and use
 594 cases.

595 For Train Ticket and Piggy Metrics, we applied end-to-end and unit testing; for SCShow-
 596 case, we applied only unit testing because it lacks a front end. Specifically, for these subjects,
 597 we executed 340, 49, and 20 test methods, respectively. Our evaluation resulted in (409 /
 598 409) passed tests, meaning that all the testing methods passed in both the *before* and *after*
 599 versions. In a way, this experiment can be characterized as *regression testing*, as it runs the
 600 same suite against both versions [83]. Table 2 summarizes our evaluation of effectiveness.

■ **Table 2** Testing Scenarios, Statistics, and Outcomes

| Use Cases | No. of Methods | Test Types | Before Monolith | After Microservice |
|---|----------------|--------------------------|-----------------|--------------------|
| Train Ticket | | | | |
| End-to-end booking flows (login, search, reserve, pay, rebook) | 20 | End-to-end (Selenium) | ✓ | ✓ |
| Administrative management features (stations, routes, prices, users) | 15 | End-to-end (Selenium) | ✓ | ✓ |
| Ticketing and travel operations (orders, payments, cancellations, notifications) | 25 | End-to-end (Selenium) | ✓ | ✓ |
| Controllers for user-facing REST endpoints (e.g., UserController, TravelController) | 120 | Unit (controller-level) | ✓ | ✓ |
| Service-layer implementations (e.g., OrderService, PriceService, TravelService) | 160 | Unit (service-layer) | ✓ | ✓ |
| Piggy Metrics | | | | |
| Front-end rendering and navigation | 1 | End-to-end (Selenium) | ✓ | ✓ |
| Account, user, and statistics controllers (API validation) | 14 | Unit (controller-level) | ✓ | ✓ |
| Core services (AccountService, NotificationService, StatisticsService) | 20 | Unit (service-layer) | ✓ | ✓ |
| Repositories for persistence (AccountRepository, UserRepository, RecipientRepository) | 9 | Unit (repository-layer) | ✓ | ✓ |
| External service clients (ExchangeRatesClient and Fallback) | 3 | Integration (API client) | ✓ | ✓ |
| Email notification functionality (EmailService) | 2 | Unit (service-layer) | ✓ | ✓ |
| SCShowcase | | | | |
| Department and Employee domain logic | 10 | Unit (service-layer) | ✓ | ✓ |
| Organization domain logic and repository | 7 | Unit (repository-layer) | ✓ | ✓ |
| REST controllers (OrganizationController) | 3 | Unit (controller-level) | ✓ | ✓ |

601 Because we created RM*—our own synthetic microservice applications—with the specific
602 goal of stress testing, they lack coherent business logic. Hence, we could not apply use-case
603 testing to these applications to assess `µjUniter`'s effectiveness. Instead, we performed a
604 two-tier structural verification. First, we conducted `ApplicationContext` loading tests using
605 the `SpringJUnit4ClassRunner` to ensure that all beans were correctly wired across the new
606 modular boundaries. Second, we performed a reachability analysis to ensure that every
607 former remote endpoint was successfully mapped to a local direct method invocation. Table 3
608 summarizes the results for these synthetic subjects.

■ **Table 3** Structural and Contextual Verification of RM* Monoliths

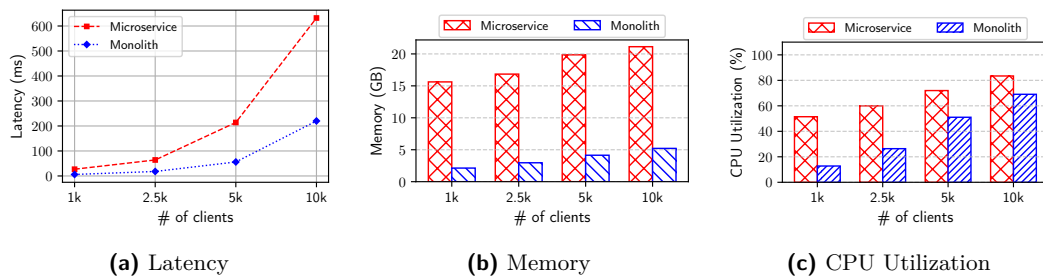
| Subjects | Compilation | Bean Wiring | Context Load | Reachability | Outcome |
|-------------|-------------|-----------------|--------------|--------------|---------|
| RM-10-0.2 | Successful | 112/112 Beans | ✓(0.8s) | 100% | Pass |
| RM-10-1.0 | Successful | 148/148 Beans | ✓(1.1s) | 100% | Pass |
| RM-100-0.02 | Successful | 1240/1240 Beans | ✓(8.4s) | 100% | Pass |

609 While our functional test suite confirms that `µjUniter` preserves the core functionality of
610 the migrated applications, these tests are insensitive to the removal of distribution-specific
611 fault handling. Because partial failure mechanisms are orthogonal to business logic, their
612 elimination should not affect observable behavior. Consequently, any residual, unreachable

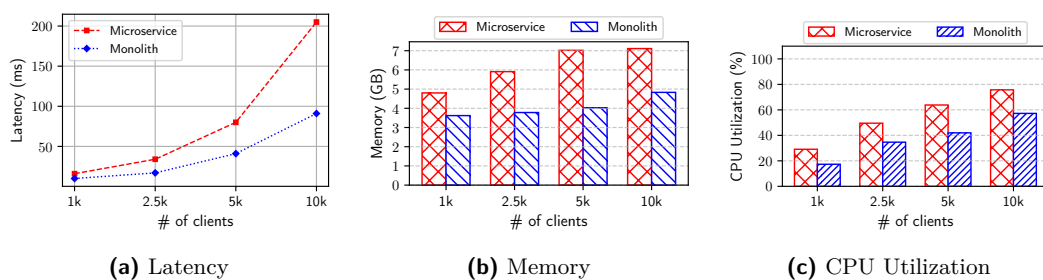
613 fault-handling code would not compromise the correctness of the application’s business logic.
 614 Instead, its removal by μ JUniter serves as a structural code simplification, reducing technical
 615 debt and streamlining execution, rather than a functional requirement. While we have not
 616 systematically evaluated μ JUniter’s dead-code removal, this omission does not impact our
 617 primary claim of functional preservation; at worst, any remaining obsolete handlers would
 618 exist as harmless, unreachable code that does not affect the application’s ability to deliver
 619 its functionality.

620 5.3 RQ2: Performance

621 We evaluate how UNIFICATION REFACTORING as realized in μ JUniter affects the performance
 622 characteristics of our subjects in terms of latency, memory usage, and CPU utilization.
 623 We measured the performance for Train Ticket and Piggy Metrics using Gatling [29], a
 624 performance-testing framework. Using Gatling, we generated workloads that simulate a
 625 specified number of clients accessing the application via the front end. The framework then
 626 reports latency, memory consumption, and CPU utilization. For subjects without front ends,
 627 we used Apache JMeter [21] to execute workloads. For the synthetic subjects, we kept the
 628 number of clients constant (30), as these subjects were used specifically for stress testing
 629 with respect to the number of microservices. Using the same workloads, we measured these
 630 characteristics for the *before* and *after* versions and compared them.



■ **Figure 4** Performance Characteristics for Train Ticket



■ **Figure 5** Performance Characteristics for Piggy Metrics

631 Figures 4, 5, and 6 depict the performance evaluation results for Train Ticket, Piggy
 632 Metrics, and SCSShowcase, respectively. Table 4 reports on the measurements for our synthetic
 633 microservice subject, RM*. Latency and resource consumption increased proportionally
 634 to the number of clients for both versions. However, the monolith versions consistently
 635 displayed lower latency and resource consumption across all workloads. These experiments
 636 confirmed the expected performance trends: replacing network communication with local calls



■ **Figure 6** Performance Characteristics for SCSShowcase

■ **Table 4** Performance Characteristics for RM* (microservice values represent the aggregate resource consumption of all distributed instances versus the single-process monolith)

| Subject | Latency (ms) | | Memory (GB) | | CPU Utilization (%) | |
|-------------|--------------|------------|--------------|-------------|---------------------|-----------|
| | Microservice | Monolith | Microservice | Monolith | Microservice | Monolith |
| RM-10-0.2 | 180 | 20 | 2.03 | 0.25 | 34 | 23 |
| RM-10-1.0 | 878 | 587 | 3.47 | 1.79 | 47 | 31 |
| RM-100-0.02 | 1594 | 281 | 37.04 | 4.94 | 71 | 52 |

637 reduces both latency and resource consumption. Specifically, across all experiments, latency
 638 decreased by $\approx 65\%$, memory consumption by $\approx 72\%$, and CPU utilization by $\approx 33\%$. The
 639 observed significant reduction in memory and CPU usage is due to μ jUniter’s elimination of
 640 both network I/O and redundant Spring Cloud infrastructure. In the monolithic version, the
 641 system no longer requires multiple embedded servlet containers, service discovery clients, and
 642 the heavy reflection-based serialization overhead associated with REST-based inter-service
 643 communication.

644 5.4 RQ3: Automation Value

645 μ jUniter automates the tedious and error-prone source code transformations required to
 646 transition from a distributed architecture to a centralized, modular monolith. To quantify
 647 this value and assess the degree of automation, we measured the lines of code removed,
 648 modified, and added using the CLOC (Count Lines of Code) tool [13], as well as the manual
 649 effort required to finalize the monolith. These results are presented in Table 5.

■ **Table 5** μ jUniter’s CLOC Transformations and Automation Ratio

| Subjects | Added (CLOC) | Removed Infrastructure | \mathcal{E}_{dist} | Modified (CLOC) | Manual (CLOC) | Automation % |
|---------------|--------------|------------------------|----------------------|-----------------|---------------|--------------|
| Train Ticket | 7520 | 2810 | 432 | 350 | 360 | 96.8% |
| Piggy Metrics | 151 | 122 | 28 | 4 | 25 | 93.2% |
| SCSShowcase | 220 | 68 | 12 | 5 | 12 | 95.1% |
| RM-10-0.2 | 289 | 92 | 12 | 34 | 0 | 100% |
| RM-10-1.0 | 558 | 110 | 25 | 153 | 0 | 100% |
| RM-100-0.02 | 3250 | 915 | 164 | 487 | 0 | 100% |

650 The Removed CLOC is categorized into two distinct groups: *Infrastructure* and \mathcal{E}_{dist} . Infra-
 651 structure removal targets Spring Cloud boilerplate, including Eureka/Consul service-discovery
 652 registrations, Zuul or Spring Cloud Gateway routing configurations, and `bootstrap.yml` files.

653 The \mathcal{E}_{dist} category represents the pruning of distribution-induced dead code identified by our
 654 *Failure Model Refiner*, including `catch` blocks for `ConnectException`, `FeignException` hand-
 655 lers, and declarative resilience infrastructure such as `@CircuitBreaker` and `@HystrixCommand`.

656 We define the Automation % as the ratio of automatically transformed logic to the
 657 total lines of added, removed, and modified code in the final monolith. While `µjUniter`
 658 achieves 100% automation for the synthetic RM* subjects, third-party applications require
 659 a small margin of manual intervention, ranging from 3% to 7%. The 100% automation for
 660 RM* is enabled by our generator, which produces services with mutually exclusive property
 661 namespaces and standardized security headers, allowing the *Failure Model Refiner* and
 662 *Inconsistency Resolver* to prune all distribution-induced artifacts and wire all beans without
 663 encountering undecidable policy conflicts.

664 In contrast, the minor human-in-the-loop requirement for third-party subjects is not
 665 a fundamental limitation of `µjUniter`, but rather a necessity for resolving high-level ar-
 666 chitectural policies that are semantically undecidable via static analysis. Specifically,
 667 manual oversight was required to merge security policies, such as consolidating multiple
 668 `WebSecurityConfigurerAdapter` beans whose different microservices defined conflicting
 669 global CORS or CSRF settings. Similarly, property namespace conflicts necessitated manual
 670 resolution when identical keys in `application.properties` were used for different purposes
 671 across separate microservice domains. Table 6 provides the breakdown of the required manual
 672 effort.

■ **Table 6** Manual Effort (CLOC / Time)

| Subject | Security Merge | Property Namespace | Total (CLOC/h) |
|---------------|-------------------|--------------------|-------------------|
| Train Ticket | 210 / 2.5h | 150 / 1.5h | 360 / 4.0h |
| Piggy Metrics | 16 / 0.3h | 9 / 0.2h | 25 / 0.5h |
| SCShowcase | 8 / 0.3h | 4 / 0.2h | 12 / 0.5h |
| Total | 234 / 3.1h | 163 / 1.9h | 397 / 5.0h |

673 Despite these manual interventions, `µjUniter` automatically handled the migration’s heavy
 674 lifting, including generating cross-module deep-copy utilities, redirecting REST clients to
 675 direct method invocations, and cleaning up signatures of checked exceptions. For a realistic
 676 third-party system like Train Ticket, the 96.8% automation rate implies that `µjUniter`
 677 successfully transformed over 11,000 lines of code, reducing the developer’s task from a
 678 massive manual refactoring project to a focused configuration review.

679 **6 Discussion**

680 In this section, we discuss the applicability of our approach, with particular emphasis on the
 681 impact of preserving modularity and the limitations of our evaluation.

682 **6.1 Applicability**

683 Conceptually, UNIFICATION REFACTORING is straightforward: extract the ASTs of distributed
 684 components, unify them into a single AST, identify the client-server endpoints, and replace
 685 network calls with local function calls, thereby producing the monolith. However, specific
 686 embodiments of UNIFICATION REFACTORING would necessarily have to adhere to the
 687 conventions of the implementation ecosystem. For example, `µjUniter`’s implementation is
 688 specifically tailored to the conventions and API of the Spring framework. UNIFICATION

689 REFACTORING should be applicable to other languages, libraries, or frameworks, but the
 690 required engineering effort can differ. Nevertheless, the main building blocks of UNIFICATION
 691 REFACTORING can be found in all major programming environments: AST manipulation and
 692 replacing network communication with local calls [15, 62]. Hence, we expect UNIFICATION
 693 REFACTORING to be widely applicable, but the availability of code analysis and transformation
 694 tools would affect the implementation effort.

6.2 Modularity and Ease of Reconciliation

696 The design of UNIFICATION REFACTORING provides significant additional benefits by pre-
 697 serving modularity during the unification process. When consolidated, originally distributed
 698 components remain isolated within individual modules. We observed that distributed deploy-
 699 ment often masks configuration heterogeneity, such as disparate Java compiler versions, which
 700 only surfaces during consolidation. Because `µjUniter` maintains modular boundaries, we
 701 were able to exploit Java’s backward compatibility: compiling each module with its original
 702 compiler while executing the unified monolith on a modern JVM. Thus, modular preservation
 703 simplifies the management of cross-version inconsistencies, which are inconsequential in
 704 isolation.

705 This structured environment is equally effective for resolving library version divergence.
 706 In a microservices architecture, independent teams may select different versions of shared
 707 libraries (e.g., JSON parsers or logging frameworks). When `µjUniter` detects such conflicts,
 708 it intentionally delegates resolution to the developer. We argue that maintaining multiple
 709 versions of the same library is frequently a symptom of postponed *technical debt* rather than
 710 a functional requirement.

711 Before performing an architectural transformation with `µjUniter`, developers should
 712 reconcile this debt. By requiring a manual update to a single, stable library version during the
 713 structural convergence phase, the developer ensures the resulting monolith is maintainable,
 714 secure, and performant. This approach avoids the complexity and overhead of runtime
 715 isolation hacks (like custom `ClassLoader`) while conforming to conventional, production-
 716 ready JVM deployment practices.

6.3 Threats to Validity and Limitations

718 **Internal Validity:** To ensure consistency and reliability of our evaluation across diverse
 719 microservice architectures, we performed targeted adjustments on our study subjects, which
 720 included the manual translation of three non-Java services in the Train Ticket benchmark
 721 and the resolution of legacy runtime errors in Piggy Metrics. While these adjustments modify
 722 the source, they were essential to ensure a uniform experimental baseline and broaden the
 723 architectural scope of our analysis. By stabilizing these subjects, we minimized the risk that
 724 runtime instability would influence our results.

725 **External Validity:** Our subject selection was guided by the objective of adhering to archi-
 726 tectural relevance and modern development standards. Although the number of evaluation
 727 subjects is focused, this choice was necessitated by our rigorous inclusion criteria: we priorit-
 728 ized high-quality Spring Java exemplars over a larger number of outdated or unmaintained
 729 repositories. While this decision may influence the immediate generalizability of the results,
 730 it ensures that our findings are grounded in stable, contemporary microservice patterns.

731 **Scale and Generalizability:** We acknowledge that large-scale enterprise systems often
 732 exhibit higher complexity than open-source benchmarks. To bridge this gap, we introduced
 733 synthetic subjects designed to simulate larger architectural footprints. Although these

734 synthetic models were calibrated for available hardware resources, they provide a valuable
735 controlled environment for observing system behavior at scale.

736 **Scope of Metrics:** Finally, this evaluation focuses on fundamental performance character-
737 istics rather than environment-specific scalability or advanced deployment metrics. Given
738 that such metrics are often highly dependent on specific infrastructure configurations, we
739 focused on providing a robust baseline analysis. We view extending this work to specialized,
740 deployment-specific performance profiling as a promising direction for future research.

741 **7 Related Work**

742 This work most closely relates to software architectural transformation and automated
743 software refactoring for distribution, which we describe in turn.

744 **Architectural Transformation**

745 As business needs and resource constraints evolve, developers often need to transform the
746 architecture of software systems. Hence, a rich body of prior work focuses on architectural
747 transformations. Common approaches to architectural transformation focus on reconstructing
748 or migrating the system structure [33, 69]. A recent state-of-the-art work investigates
749 three frameworks and four case studies for building modular monoliths as an alternative
750 or precursor to microservices [76]. Another recent work evaluates the performance and
751 migration effort required to transform a monolith into microservices, creating a centralized
752 service as an intermediate step [19]. IBM’s Mono2Micro is a practical tool that decomposes
753 monolithic Java applications into microservices by spatio-temporally clustering static and
754 runtime data [39]. Micro2Micro optimizes the service boundaries of microservice systems and
755 creates a centralized version as an intermediate step [12]. Work on service modularization,
756 such as Service Cutter, analyzes coupling criteria to guide architectural slicing into well-
757 defined services [31]. Google’s Service Weaver provides a framework for writing modular
758 monolith applications in Go that can be deployed either as a single process or as distributed
759 microservices, preserving modularity while enabling runtime distribution [30], albeit requiring
760 that applications be built from the ground up using a specific programming model. More
761 recently, AI-driven approaches leverage LLMs to detect module boundaries and determine
762 transformations that improve various characteristics of Java enterprise systems [84, 65].

763 These works share similar objectives—transforming architectures to optimize system
764 performance. However, their transformations and tools focus on either increasing distribution
765 by splitting monoliths into microservices or on optimizing existing microservice systems,
766 rather than on reducing distribution. Although tools that create a centralized system version
767 as an intermediate step are similar to μ JUniter, we codify this architectural transformation
768 as a new source-to-source code refactoring—UNIFICATION REFACTORING—and demonstrate
769 that it can be largely automated. To that end, this work has identified how to consolidate
770 microservices into a modular monolith through automated refactoring, preserving modularity
771 while eliminating network communication overhead. Unlike architectural reconstruction tools,
772 μ JUniter’s pipeline merges microservice ASTs and rewires endpoints across modules. To our
773 knowledge, this work is the first to identify the need for improved support for migrating
774 microservices to a monolith to reduce distribution.

775 Automated Refactoring for Distribution

776 Several prior works present automated tools that assist with architectural refactoring in
 777 distributed systems. An automatic framework analyzes Java systems, reconstructs their
 778 conceptual architecture, and refactors them towards that architecture [73]. Another study
 779 critically reviews and categorizes several automatic refactoring strategies for decomposing
 780 monoliths into microservices [26]. Log2MS extracts execution logs to drive automated refact-
 781 oring of monoliths into microservices, optimizing partitions based on runtime behavior [48].
 782 Other automated refactoring tools, such as ARIES [6], apply rule-based source-to-source
 783 transformations to restructure legacy Java systems, while MicroDepGraph [70] builds de-
 784 pendency graphs from Java code and deployment descriptors to recommend microservice
 785 extractions. D-Goldilocks automatically rearchitects distributed systems to achieve an op-
 786 timal degree of distribution [4]. As an intermediate step, D-Goldilocks also introduces a
 787 unification process, albeit specific to Node.js JavaScript applications, applied to running
 788 systems, and requiring programmer annotations.

789 These automated approaches also either focus on partitioning monoliths or redistributing
 790 distributed systems. None of them focuses on the problem of reducing distribution as an
 791 automated refactoring. Although we draw on prior automated refactoring solutions, our
 792 program analysis and transformation routines focus specifically on consolidating distributed
 793 microservices into a monolith.

794 **8** Conclusion

795 This paper presents a novel refactoring—UNIFICATION REFACTORING—that transforms a
 796 system comprising distributed components into a semantically equivalent centralized version,
 797 thereby closing the loop on architecture migration. Its source-to-source transformation
 798 pipeline merges abstract syntax trees and bridges control flow, consolidating distributed
 799 components into a centralized system. We reify UNIFICATION REFACTORING as μ jUniter,
 800 an automated refactoring framework that migrates Java Spring Boot microservices into
 801 a functionally equivalent modular monolith. Our evaluation of μ jUniter shows that the
 802 resulting monolith maintains functionality while reducing latency and resource utilization.
 803 μ jUniter’s additional benefits include systematically refining failure models and pruning
 804 distribution-induced dead code, ensuring the removal of redundant network-handling logic in
 805 the resulting monolith. Overall, μ jUniter provides substantial automation by eliminating the
 806 need to manually transform microservice source code and automatically handling over 96%
 807 of refactoring tasks. As technological advancements continue to impact software architecture,
 808 the need for automated architectural migrations will only grow. UNIFICATION REFACTORING
 809 and μ jUniter provide innovative designs and insights for this evolving field. As future work, we
 810 plan to extend μ jUniter’s applicability to other languages and frameworks. Another direction
 811 is to extend μ jUniter’s functionality to support the redistribution of microservices, with
 812 unification as an intermediate step. Because modern systems often need to both increase and
 813 decrease distribution to accommodate evolving requirements and constraints, UNIFICATION
 814 REFACTORING fills the gap in support for architectural transformation.

815 — References —

- 816 1 Chaima Abid, Vahid Alizadeh, Marouane Kessentini, Thiago do Nascimento Ferreira, and
 817 Danny Dig. 30 years of software refactoring research: A systematic literature review. *arXiv*
 818 *preprint arXiv:2007.02194*, 2020.

- 819 2 Chaima Abid, Marouane Kessentini, Vahid Alizadeh, Mouna Dhaouadi, and Rick Kazman.
820 How does refactoring impact security when improving quality? a security-aware refactoring
821 approach. *IEEE Transactions on Software Engineering*, 48(3):864–878, 2020.
- 822 3 Omar Al-Debagy and Peter Martinek. A comparative review of microservices and monolithic
823 architectures. In *2018 IEEE 18th International Symposium on Computational Intelligence
824 and Informatics (CINTI)*, pages 000149–000154. IEEE, 2018.
- 825 4 Kijin An and Eli Tilevich. D-goldilocks: Automatic redistribution of remote functionalities for
826 performance and efficiency. In *2020 IEEE 27th International Conference on Software Analysis,
827 Evolution and Reengineering (SANER)*, pages 251–260. IEEE, 2020.
- 828 5 Lujo Bauer, Andrew W Appel, and Edward W Felten. Mechanisms for secure modular
829 programming in Java. *Software: Practice and Experience*, 33(5):461–480, 2003.
- 830 6 Gabriele Bavota, Andrea De Lucia, Andrian Marcus, Rocco Oliveto, and Fabio Palomba.
831 Supporting extract class refactoring in eclipse: The aries project. In *2012 34th International
832 Conference on Software Engineering (ICSE)*, pages 1419–1422. IEEE, 2012.
- 833 7 Ankur Bawiskar, Prashant Sawant, Vinayak Kankate, and BB Meshram. Spring framework: a
834 companion to JavaEE. *IJCEM*, 1:41–49, 2012.
- 835 8 Thomas Betts. To microservices and back again – why segment went back to a Monolith, Apr
836 2020. URL: <https://www.infoq.com/news/2020/04/microservices-back-again/>.
- 837 9 Anshita Bhasin. Exploring Amazon Prime Video’s architecture: Migrating from microservices
838 to monolith for... Medium, May 2024. URL: <https://medium.com/@anshita.bhasin/exploring-amazon-prime-videos-architecture-migrating-from-microservices-to-monolith-for-aacbf9fab373>.
- 841 10 Craig Box. Introducing istiod: simplifying the control plane. Technical report, Istio, 2020.
842 URL: <https://istio.io/v1.5/blog/2020/istiod/>.
- 843 11 Tomas Cerny, Md Showkat Hossain Chy, Muhmmad Ashfakur Rahman Arju, Korn Sooksatra,
844 Amr S Abdelfattah, and Valentina Lenarduzzi. A multi-variant benchmark for microservice
845 systems in software engineering research. In *European Conference on Software Architecture*,
846 pages 21–29. Springer, 2024.
- 847 12 Md Showkat Hossain Chy, Korn Sooksatra, Jorge Yero, and Tomas Cerny. Benchmarking
848 Micro2Micro transformation: an approach with GNN and VAE. *Cluster Computing*, 27(4):4171–
849 4185, 2024.
- 850 13 Al Daniaal. Cloc: Count lines of code, 2017. URL: <https://cloc.sourceforge.net/>.
- 851 14 Paul Deitel. Understanding Java 9 modules. Oracle, 2017. URL: <https://www.oracle.com/corporate/features/understanding-java-9-modules.html>.
- 853 15 DevZery. Asts meaning: A complete programming guide. <https://www.devzery.com/post/ast-meaning-a-complete-guide-to-abstract-syntax-trees-in-programming>, 2025.
- 855 16 Edsger Wybe Dijkstra. Structured programming. 1970.
- 856 17 Oliver Drotbohm. Spring blog. Introducing Spring Modulith, Oct 2022. URL: <https://spring.io/blog/2022/10/21/introducing-spring-modulith>.
- 858 18 Hassan Farsi, Driss Allaki, Abdeslam En-Nouaary, and Mohamed Dahchour. A graph-based
859 solution to deal with cyclic dependencies in microservices architecture. In *2022 9th International
860 Conference on Future Internet of Things and Cloud (FiCloud)*, pages 254–259. IEEE, 2022.
- 861 19 Diogo Faustino, Nuno Gonçalves, Manuel Portela, and António Rito Silva. Stepwise migration
862 of a monolith to a microservice architecture: Performance and migration effort evaluation.
863 *Performance Evaluation*, 164:102411, 2024.
- 864 20 Stefan Fischer, Pirmin Urbanke, Rudolf Ramler, Monika Steidl, and Michael Felderer. An
865 overview of microservice-based systems used for evaluation in testing and monitoring: A
866 systematic mapping study. In *Proceedings of the 5th ACM/IEEE International Conference on
867 Automation of Software Test (AST 2024)*, pages 182–192, 2024.
- 868 21 The Apache Software Foundation. Apache JMeter™, 1998. URL: <https://jmeter.apache.org/>.
- 869

- 870 22 The Apache Software Foundation. Apache Maven project. <https://maven.apache.org/what-is-maven.html>, 2004.
- 871
- 872 23 The Apache Software Foundation. Guide to working with multiple modules, 2004. URL: <https://maven.apache.org/guides/mini/guide-multiple-modules.html>.
- 873
- 874 24 Martin Fowler. *Refactoring: improving the design of existing code*. Addison-Wesley Professional, 2018.
- 875
- 876 25 Jonas Fritzsich, Justus Bogner, Stefan Wagner, and Alfred Zimmermann. Microservices migration in industry: intentions, strategies, and challenges. In *2019 IEEE International Conference on Software Maintenance and Evolution (ICSME)*, pages 481–490. IEEE, 2019.
- 877
- 878
- 879 26 Jonas Fritzsich, Justus Bogner, Alfred Zimmermann, and Stefan Wagner. From monolith to microservices: A classification of refactoring approaches. In *International Workshop on Software Engineering Aspects of Continuous Development and New Paradigms of Software Production and Deployment*, pages 128–141. Springer, 2018.
- 880
- 881
- 882
- 883 27 FudanSELab. Train ticket a benchmark microservice system, August 2022. URL: <https://github.com/FudanSELab/train-ticket>.
- 884
- 885 28 John Fx. Serverless was a big mistake... says amazon. YouTube, May 2023. URL: <https://www.youtube.com/watch?v=qQk94CjRvIs>.
- 886
- 887 29 Gatling. Gatling documentation. <https://docs.gatling.io/>, 2012.
- 888
- 889 30 Sanjay Ghemawat, Robert Grandl, Srdjan Petrovic, Michael Whittaker, Parveen Patel, Ivan Posva, and Amin Vahdat. Towards modern development of cloud applications. In *Proceedings of the 19th Workshop on Hot Topics in Operating Systems*, pages 110–117, 2023.
- 890
- 891 31 Michael Gysel, Lukas Kölbener, Wolfgang Giersche, and Olaf Zimmermann. Service cutter: A systematic approach to service decomposition. In *European Conference on Service-Oriented and Cloud Computing*, pages 185–200. Springer, 2016.
- 892
- 893
- 894 32 Muhammad Hafiz Hasan, Mohd Hafeez Osman, Novia Indriaty Admodisastro, and Muhammad Sufri Muhammad. A quality driven framework for decomposing legacy monolith applications to microservice architecture. Research Square, 2023. doi:10.21203/rs.3.rs-3060410/v1.
- 895
- 896
- 897
- 898 33 Wilhelm Hasselbring. Software architecture: Past, present, future. In *The essence of software engineering*, pages 169–184. Springer International Publishing Cham, 2018.
- 899
- 900 34 Amy E Hodler and Mark Needham. Graph data science using neo4j. In *Massive Graph Analytics*, pages 433–457. Chapman and Hall/CRC, 2022.
- 901
- 902 35 Scott M. Fulton III. Amazon prime video’s microservices move doesn’t lead to a monolith after all. The New Stack, Jun 2023. URL: <https://thenewstack.io/amazon-prime-video-s-microservices-move-doesnt-lead-to-a-monolith-after-all/>.
- 903
- 904
- 905 36 ISO/IEC/IEEE. Software engineering — software life cycle processes — maintenance, 2022. doi:10.1109/IEEESTD.2022.9690131.
- 906
- 907 37 JetBrains. JetBrains Microservices. <https://www.jetbrains.com/lp/devecosystem-2022/microservices/>, 2023.
- 908
- 909 38 Xinyang Jia. The role and importance of software testing in software quality management. *Journal of Industry and Engineering Management*, 1(4):39–44, 2023.
- 910
- 911 39 Anup K Kalia, Jin Xiao, Rahul Krishna, Saurabh Sinha, Maja Vukovic, and Debasish Banerjee. Mono2micro: a practical and effective tool for decomposing monolithic Java applications to microservices. In *Proceedings of the 29th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, pages 1214–1224, 2021.
- 912
- 913
- 914
- 915
- 916 40 Raffi Khatchadourian, Jason Sawin, and Atanas Rountev. Automated refactoring of legacy Java software to enumerated types. In *2007 IEEE International Conference on Software Maintenance*, pages 224–233. IEEE, 2007.
- 917
- 918
- 919 41 Marcin Kolny. Scaling up the prime video audio/video monitoring service and reducing costs by 90%, Mar 2023. URL: <https://www.wudsn.com/productions/www/site/news/2023/2023-05-08-microservices-01.pdf>.
- 920
- 921

- 922 42 Koushik Kothagal. *Modular Programming in Java 9: Build large scale applications using Java*
923 *modularity and Project Jigsaw*. Packt Publishing Ltd, 2017.
- 924 43 Anthony Kwan, Hans-Arno Jacobsen, Allen Chan, and Suzette Samoojh. Microservices in
925 the modern software world. In *Proceedings of the 26th Annual International Conference on*
926 *Computer Science and Software Engineering*, pages 297–299, 2016.
- 927 44 Young-Woo Kwon. Automated s/w reengineering for fault-tolerant and energy-efficient dis-
928 tributed execution. In *2013 IEEE International Conference on Software Maintenance*, pages
929 582–585. IEEE, 2013.
- 930 45 Young-Woo Kwon and Eli Tilevich. Cloud refactoring: automated transitioning to cloud-based
931 services. *Automated Software Engineering*, 21(3):345–372, 2014.
- 932 46 Mahesh Lal. *Neo4j graph data modeling*. Packt Publishing Ltd, 2015.
- 933 47 Duc Minh Le, Pooyan Behnamghader, Joshua Garcia, Daniel Link, Arman Shahbazian, and
934 Nenad Medvidovic. An empirical study of architectural change in open-source software systems.
935 In *2015 IEEE/ACM 12th working conference on mining software repositories*, pages 235–245.
936 IEEE, 2015.
- 937 48 Bo Liu, Jingliu Xiong, Qiorong Ren, Shmuel Tyszberowicz, and Zheng Yang. Log2ms: a
938 framework for automated refactoring monolith into microservices using execution logs. In
939 *2022 IEEE International Conference on Web Services (ICWS)*, pages 391–396. IEEE, 2022.
- 940 49 Yin Liu, Kijin An, and Eli Tilevich. Rt-trust: automated refactoring for trusted execution under
941 real-time constraints. In *Proceedings of the 17th ACM SIGPLAN International Conference on*
942 *Generative Programming: Concepts and Experiences*, pages 175–187, 2018.
- 943 50 Alexander Lukyanchikov. Piggy metrics, July 2018. URL: <https://github.com/sqshq/piggy-metrics>.
- 944
- 945 51 Nabor C Mendonça, Craig Box, Costin Manolache, and Louis Ryan. The monolith strikes
946 back: Why istio migrated from microservices to a monolithic architecture. *IEEE software*,
947 38(5):17–22, 2021.
- 948 52 Naftaly H Minsky and Victoria Ungureanu. Law-governed interaction: a coordination and
949 control mechanism for heterogeneous distributed systems. *ACM Transactions on Software*
950 *Engineering and Methodology (TOSEM)*, 9(3):273–305, 2000.
- 951 53 Mayank Mishra, Shruti Kunde, and Manoj Nambiar. Cracking the monolith: Challenges in
952 data transitioning to cloud native architectures. In *Proceedings of the 12th European conference*
953 *on software architecture: companion proceedings*, pages 1–4, 2018.
- 954 54 Piotr Mińkowski. Microservices with Spring cloud advanced demo project, May 2024. URL:
955 <https://github.com/piomin/sample-spring-microservices-new/tree/master>.
- 956 55 Ragnar Mogk, Lars Baumgärtner, Guido Salvaneschi, Bernd Freisleben, and Mira Mezini.
957 Fault-tolerant distributed reactive programming. In *32nd European Conference on Object-*
958 *Oriented Programming (ECOOP 2018)*, pages 1–1. Schloss Dagstuhl–Leibniz-Zentrum für
959 Informatik, 2018.
- 960 56 Rodrigo Morales, Rubén Saborido, Foutse Khomh, Francisco Chicano, and Giuliano Antoniol.
961 Earmo: An energy-aware refactoring approach for mobile apps. In *Proceedings of the 40th*
962 *International Conference on Software Engineering*, pages 59–59, 2018.
- 963 57 Conor Muldoon, Levent Görgü, John J O’Sullivan, Wim G Meijer, Bartholomew Masterson,
964 and Gregory MP O’Hare. Engineering testable and maintainable software with Spring Boot
965 and React. *Authorea Preprints*, 2023.
- 966 58 Kshirasagar Naik and Priyadarshi Tripathy. *Software testing and quality assurance: theory*
967 *and practice*. John Wiley & Sons, 2011.
- 968 59 Iulian Neamtiu, Jeffrey S Foster, and Michael Hicks. Understanding source code evolution
969 using abstract syntax tree matching. In *Proceedings of the 2005 international workshop on*
970 *Mining software repositories*, pages 1–5, 2005.
- 971 60 Vikram Nitin, Shubhi Asthana, Baishakhi Ray, and Rahul Krishna. Cargo: AI-guided
972 dependency analysis for migrating monolithic applications to microservices architecture.

- 973 In *Proceedings of the 37th IEEE/ACM International Conference on Automated Software*
974 *Engineering*, pages 1–12, 2022.
- 975 **61** Rory V O'Connor, Peter Elger, and Paul M Clarke. Continuous software engineering—a mi-
976 croservices architecture perspective. *Journal of Software: Evolution and Process*, 29(11):e1866,
977 2017.
- 978 **62** Knowledgelark Encyclopedia of Daily Life. Handling network request mocks in tests. <https://www.knowledgelark.com/archives/5263>, 2025.
- 980 **63** William F Opdyke. *Refactoring object-oriented frameworks*. University of Illinois at Urbana-
981 Champaign, 1992.
- 982 **64** Amey Padvekar and Vikaskumar Gupta. Comparative analysis of monolithic vs. distributed
983 architecture. *International Journal of Advanced Research in Science, Communication and*
984 *Technology*, pages 433–442, 06 2024. doi:10.48175/IJARSCT-18946.
- 985 **65** Indranil Palit and Tushar Sharma. Generating refactored code accurately using reinforcement
986 learning. *arXiv preprint arXiv:2412.18035*, 2024.
- 987 **66** Renaud Pawlak, Martin Monperrus, Nicolas Petitprez, Carlos Noguera, and Lionel Seinturier.
988 Spoon: A library for implementing analyses and transformations of Java source code. *Software:*
989 *Practice and Experience*, 46(9):1155–1179, 2016.
- 990 **67** Perforce, Inc. 2023 Java developer productivity report. [https://www.jrebel.com/resources](https://www.jrebel.com/resources/java-developer-productivity-report-2023)
991 [/java-developer-productivity-report-2023](https://www.jrebel.com/resources/java-developer-productivity-report-2023), 2023.
- 992 **68** Evaggelia Pitoura and Bharat Bhargava. Data consistency in intermittently connected
993 distributed systems. *IEEE Transactions on knowledge and data engineering*, 11(6):896–915,
994 2002.
- 995 **69** Konstantinos Plakidas, Daniel Schall, and Uwe Zdun. Software migration and architecture
996 evolution with industrial platforms: A multi-case study. In *European Conference on Software*
997 *Architecture*, pages 336–343. Springer, 2018.
- 998 **70** Mohammad Imranur Rahman, Sebastiano Panichella, and Davide Taibi. A curated dataset of
999 microservices-based systems. In *Joint of the Summer School on Software Maintenance and*
1000 *Evolution*. CEUR-WS, 2019.
- 1001 **71** Mark Russinovich. Microservices: An application revolution powered by the cloud, May 2023.
1002 URL: [https://azure.microsoft.com/en-us/blog/microservices-an-application-revol](https://azure.microsoft.com/en-us/blog/microservices-an-application-revolution-powered-by-the-cloud/)
1003 [ution-powered-by-the-cloud/](https://azure.microsoft.com/en-us/blog/microservices-an-application-revolution-powered-by-the-cloud/).
- 1004 **72** Max Schäfer, Torbjörn Ekman, and Oege De Moor. Sound and extensible renaming for Java.
1005 In *Proceedings of the 23rd ACM SIGPLAN conference on Object-oriented programming systems*
1006 *languages and applications*, pages 277–294, 2008.
- 1007 **73** Frederik Schmidt, Stephen G MacDonell, and Andrew M Connor. An automatic architecture
1008 reconstruction and refactoring framework. In *Software engineering research, management and*
1009 *applications 2011*, pages 95–111. Springer, 2012.
- 1010 **74** Sheldon Smith, Ethan Robinson, Timmy Frederiksen, Trae Stevens, Tomas Cerny, Miroslav
1011 Bures, and Davide Taibi. Benchmarks for end-to-end microservices testing. In *2023 IEEE*
1012 *International Conference on Service-Oriented System Engineering (SOSE)*, pages 60–66. IEEE,
1013 2023.
- 1014 **75** Jacopo Soldani, Damian Andrew Tamburri, and Willem-Jan Van Den Heuvel. The pains and
1015 gains of microservices: A systematic grey literature review. *Journal of Systems and Software*,
1016 146:215–232, 2018.
- 1017 **76** Ruoyu Su and Xiaozhou Li. Modular monolith: Is this the trend in software architecture? In
1018 *Proceedings of the 1st International Workshop on New Trends in Software Architecture*, pages
1019 10–13, 2024.
- 1020 **77** Kridanto Surendro, Wikan Danar Sunindyo, et al. Circuit breaker in microservices: State of
1021 the art and future prospects. In *IOP Conference Series: Materials Science and Engineering*,
1022 volume 1077, page 012065. IOP Publishing, 2021.

- 1023 **78** Wesley Tansey and Eli Tilevich. Annotation refactoring: inferring upgrade transformations for
1024 legacy applications. In *Proceedings of the 23rd ACM SIGPLAN conference on Object-oriented*
1025 *programming systems languages and applications*, pages 295–312, 2008.
- 1026 **79** Wesley Tansey and Eli Tilevich. Efficient automated marshaling of C++ data structures
1027 for MPI applications. In *2008 IEEE International Symposium on Parallel and Distributed*
1028 *Processing*, pages 1–12, 2008. doi:10.1109/IPDPS.2008.4536307.
- 1029 **80** Eli Tilevich and Yannis Smaragdakis. J-Orchestra: Automatic Java application partitioning.
1030 In Boris Magnusson, editor, *ECOOP 2002 — Object-Oriented Programming*, pages 178–204,
1031 Berlin, Heidelberg, 2002. Springer Berlin Heidelberg.
- 1032 **81** Victor Velepucha and Pamela Flores. A survey on microservices architecture: Principles,
1033 patterns and migration challenges. *IEEE access*, 11:88339–88358, 2023.
- 1034 **82** Andrew Walker, Ian Laird, and Tomas Cerny. On automatic software architecture reconstruc-
1035 tion of microservice applications. In *Information Science and Applications: Proceedings of*
1036 *ICISA 2020*, pages 223–234. Springer, 2021.
- 1037 **83** W Eric Wong, Joseph R Horgan, Saul London, and Hiralal Agrawal. A study of effective
1038 regression testing in practice. In *PROCEEDINGS The Eighth International Symposium On*
1039 *Software Reliability Engineering*, pages 264–274. IEEE, 1997.
- 1040 **84** Yisen Xu, Feng Lin, Jinqiu Yang, Nikolaos Tsantalis, et al. Mantra: Enhancing automated
1041 method-level refactoring with contextual rag and multi-agent llm collaboration. *arXiv preprint*
1042 *arXiv:2503.14340*, 2025.
- 1043 **85** Yung-Yu Zhuang, Wei Kuo, and Shang-Chun Tseng. Resolving the Java representation exposure
1044 problem with an AST-based deep copy and flexible alias ownership system. *Electronics*,
1045 13(2):350, 2024.
- 1046 **86** Mohamed Zouari and Ismael Bouassida Rodriguez. Towards automated deployment of
1047 distributed adaptation systems. In *European Conference on Software Architecture*, pages
1048 336–339. Springer, 2013.