

Babel: A Framework for Developing Performant and Dependable Distributed Protocols

Pedro Fouto, Pedro Ákos Costa, Nuno Preguiça, João Leitão
NOVA LINCS & NOVA FCT

{p.fouto,pah.costa}@campus.fct.unl.pt, {nuno.preguiça,jc.leitao}@fct.unl.pt

Abstract—Prototyping and implementing distributed algorithms, particularly those that address challenges related with fault-tolerance and dependability, is a time consuming task. This is, in part, due to the need of addressing low level aspects such as management of communication channels, controlling timeouts or periodic tasks, and dealing with concurrency issues. This has a significant impact for researchers that want to build prototypes for conducting experimental evaluation; practitioners that want to compare different design alternatives/solutions; and even for practical teaching activities on distributed algorithms courses.

In this paper we present Babel, a novel framework to develop, implement, and execute distributed protocols and systems. Babel promotes an event driven programming and execution model that simplifies the task of translating typical specifications or descriptions of algorithms into performant prototypes, while allowing the programmer to focus on the relevant challenges of these algorithms by transparently handling time consuming low level aspects. Furthermore, Babel provides, and allows the definition of, networking components that can capture different network capabilities (e.g., P2P, Client/Server, φ -accrual Failure Detector), making the code mostly independent from the underlying communication aspects. Babel was built to be generic and can be used to implement a wide variety of different classes of distributed protocols. We conduct our experimental work with two relevant case studies, a Peer-to-Peer application and a State Machine Replication application, that show the generality and ease of use of Babel and present competitive performance when compared with significantly more complex implementations.

Index Terms—Distributed Protocols, Framework, P2P, Consensus

I. INTRODUCTION

Distributed systems are being increasingly used to support distinct activities in our everyday life. With the transition to digital platforms, distributed systems continue to grow in scale and consequently in complexity, with strict fault-tolerance and dependability requirements. This, in turn, increases the pressure for developing, validating, and testing different alternatives for building dependable abstractions that support the operation of many of these systems.

The pressure to develop better and more performant distributed systems often requires comparing different alternatives found in the literature. This entails creating concrete and correct implementations of solutions described in research papers

This work was partially supported by Fundação para a Ciência e Tecnologia (FCT) under the grants PTDC/CCI-INF/32038/2017 (NG-STORAGE), UIDB/04516/2020 (NOVA LINCS), SFRH/BD/143668/2019 and SFRH/BD/144023/2019. Experiments presented in this paper were carried out using the Grid'5000 testbed, supported by a scientific interest group hosted by Inria and including CNRS, RENATER and several Universities as well as other organizations (see <https://www.grid5000.fr/>).

and books [1]–[3] for comparing their performance under different conditions. This is the case for researchers developing novel solutions that wish to compare their proposals with the existing state of the art; for practitioners, that want to make informed decisions regarding the design and implementation of systems under development and operation; and even in classrooms, where students being exposed to dependable algorithms or large-scale systems can significantly benefit from the opportunity to implement and operate practical solutions [4].

Many (fault-tolerant) distributed algorithms/protocols are conceptually simple to describe, with relevant examples including agreement protocols [5]–[7], common abstractions such as state-machine replication [8], large-scale peer-to-peer protocols [9]–[11], among others. Implementing these protocols tends to be error-prone and time consuming [12], negatively affecting the activities of researchers, practitioners, and professors in the fault-tolerant and dependability community. This is not so much due to the complexity of the algorithms themselves, but due to the fact that such implementations require handling several low level aspects such as the use of communication primitives, scheduling and handling timed/periodic actions, or even deal with non-trivial concurrency issues that naturally arise when implementing complex systems. Additionally, stand-alone implementations require a large effort in lines of code, leading to the logic of protocols, usually described in pseudo-code in research papers and text books [4], to be lost within the implementation details.

An alternative to practical implementation is to take advantage of simulators, such as Peersim [13] or NS2/NS3 [14], [15]. However, there are several arguments that can be made against these approaches: *i*) such implementations can lead to implementations that incorrectly capture a realistic execution environment (e.g., in Peersim it is extremely easy for a node to access the internal state of another node in the system without requiring any communication step); *ii*) simulators make it hard to access the performance of competing alternatives which is many times essential for researchers and practitioners alike; *iii*) some simulators might require an amount of implementation effort that is close to a real implementation (e.g., NS2 was famous for requiring very complex protocol implementations); and *iv*) simulators required implementing the same protocol twice (the simulator code and the real code), without any guarantees that the real code matches the specification of the simulator code.

Actor-based frameworks, such as Akka [16] or Erlang-OTP [17], offer distributed programming abstractions that allow experienced system engineers to implement distributed protocols quickly and efficiently. However, these high level abstractions sometimes make it difficult to have the necessary fine-grained control over the execution, which is fundamental for high performant implementations. The adoption of a functional programming model also requires additional effort in translating the specifications, usually presented in an imperative way in the literature.

In this paper we propose Babel, a novel Java framework that provides a programming environment for implementing distributed protocols and systems. Babel provides an event driven programming and execution model that simplifies the task of translating typical specifications and descriptions of algorithms in practical prototypes with no significant performance bottlenecks. As we discuss further ahead, the abstractions provided by Babel are sufficiently generic to enable its usage for implementing a wide variety of distributed protocols and abstractions. Furthermore, Babel allows protocols to be implemented independently of the networking abstractions used to support the communication between different processes, and shields the programmer from dealing with potentially complex concurrency aspects within, and across, distributed protocols. This does not mean that Babel is only useful for prototyping, or that it is non-performant. On the contrary, as we show in our evaluation, Babel was implemented with performance in mind, allowing developers to implement efficient, production-ready protocols. The source code for Babel is available at <https://github.com/pfouto/babel-core>.

The remainder of this paper is organized as follows: we discuss relevant related work in Section II. In Section III, we discuss the features and APIs of Babel. Section IV presents our experimental work with two relevant case studies: a Peer-to-Peer application and a State Machine Replication application using MultiPaxos [18]. Finally, we conclude the paper in Section V by providing some pointers for future work.

II. RELATED WORK

We highlight some works that share the goals of Babel of providing quick development and reuse of practical distributed protocols and systems.

Appia [19] is a toolkit that takes advantage of the Java inheritance model to allow developers to create their own protocols by extending base elements of Appia that control the execution, which our own work also leverages. Appia introduces the concepts of channels and sessions, which provide developers with more control over the binding of protocols to create different types of services within a single application however, at the cost of requiring developers to stack protocols, which limits interactions among these different protocols. Furthermore, Appia has a single execution thread for executing all protocols that compose the distributed system, which incurs in non-negligible performance bottlenecks.

Cactus [20] is a C++ framework aiming at extracting more performance from a system and being more expressive. To

this end, Cactus provides meta-protocols that can be stacked together and that can be composed of various smaller protocols which can be executed asynchronously. Unfortunately, Cactus does not offer any type of concurrency management, putting that burden on the developer.

Yggdrasil [21] is a C framework that aims to be lightweight and efficient. It was designed to allow the execution of distributed applications in wireless ad-hoc scenarios, and has extensions for traditional wired IP networks. Yggdrasil provides a clean development interface focused on the processing of four types of events: *i*) network messages; *ii*) timed events; *iii*) protocol indirect interactions; and *iv*) protocol direct interactions. The last two events are used to implement message passing, allowing for the concurrent execution of protocols while shielding the developer from complex concurrency issues. However, Yggdrasil only supports a single active network abstraction, meaning that it can only execute either protocols that expect wireless ad hoc or wired IP networking in each single Yggdrasil instance. Furthermore, being tailored for C development, it requires the developer to be technically disciplined as to not incur in frequent invalid memory accesses, which can lead to time consuming development.

It is also important to note that some works have proposed Java frameworks tailored for distributed systems, but focus on the understanding of distributed protocols. This line of work includes ViSiDia [22] and the work proposed in [23]. Other works have focused on the experimental evaluation and performance assessment of distributed systems in realistic settings including SPLAY [4] and Kollaps [24]. These works are complementary to our own. Whereas we focus in providing adequate tooling for the quick, correct, and performant implementation of distributed protocols and systems, these works simplify the task of conducting performance assessment of such implementations.

Simulators, such as Peersim [13] or NS2/NS3 [14], [15], are an alternative for quickly developing and testing distributed systems. However, simulators often fail to completely model a realistic execution environment and may even allow incorrect implementations (e.g. implementations where nodes access the internal state of other nodes). More realistic simulators tend to require an implementation effort that is close to a real implementation (e.g., NS2 was famous for requiring very complex protocol implementations). In contrast, Babel is not a simulator, but rather a framework that simplifies the implementation of distributed systems that can be tested and deployed on real hardware.

III. BABEL

Babel is a framework that aims to simplify the development of distributed protocols within a process that executes in real hardware. A process can execute any number of (different) protocols that communicate with each other or/and protocols in different processes. Babel simplifies the development by enabling the developer to focus on the logic of the protocol, without having to deal with low level complexities associated

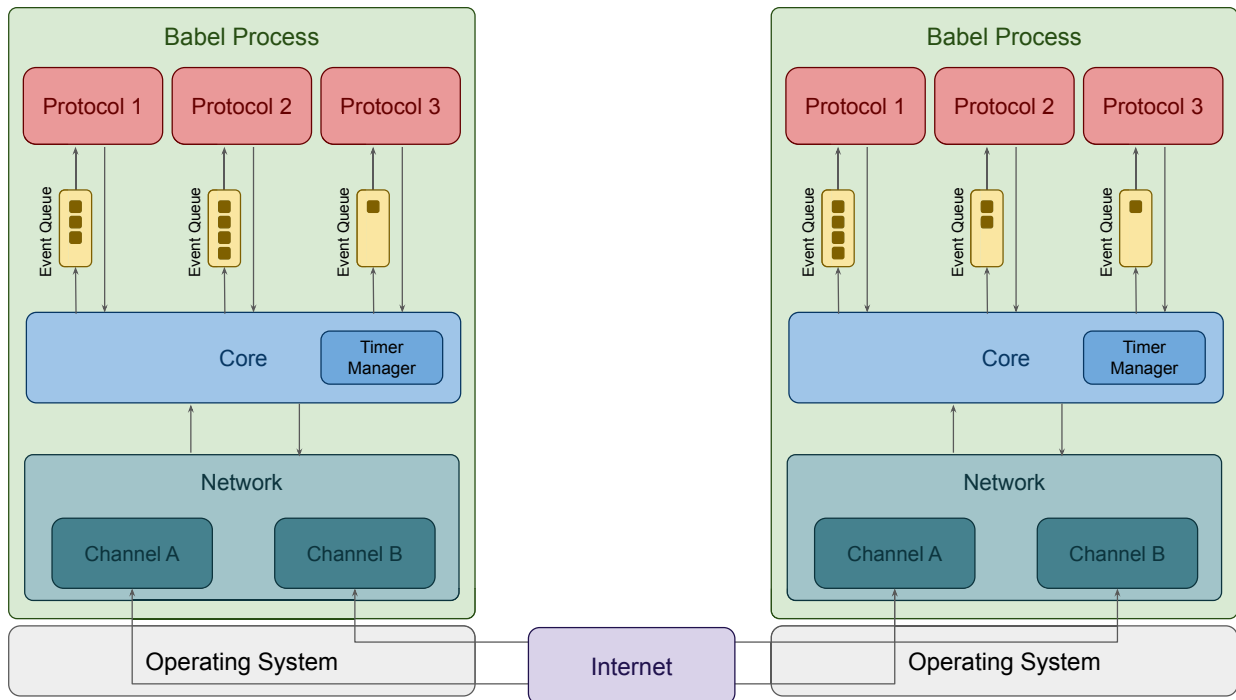


Fig. 1: Architecture of Babel

with typical distributed systems implementations. These complexities include interactions among (local) protocols, handling message passing and communication aspects, handling timers, and concurrency-control aspects within, and across, protocols (while enabling different protocols within a process to progress independently). Notably, Babel hides communication complexities behind abstractions called *channels* that can be extended/modified by the developer, with Babel already offering several alternatives that capture different capabilities (e.g., P2P, Client/Server, φ -accrual Failure Detector). Babel is implemented in Java, taking advantage of its inheritance mechanisms, such that developers extend abstract classes provided by the framework to develop their own protocols and solutions. The strong typing provided by Java allows the framework to easily enforce expected behavior, while at the same time offering enough flexibility for the developer to implement any type of distributed protocol or system.

Figure 1 presents the architecture of Babel. In the example, there are two processes executing Babel, each process being composed by three protocols and two network channels for inter-process communication. Naturally, any distributed system operating in the real world will be composed by more than two processes. The Babel framework is composed by three main components, which we now detail:

1) *Protocols*: Protocols are implemented by developers (i.e., the users of the Babel framework), and encode all the behavior of the distributed system being designed. Each protocol is modeled as a state machine whose state evolved by the

reception and processing of (external) events. For this purpose, each protocol contains an event queue from which events are retrieved. These events can be *Timers*, *Channel Notifications* from the network layer, *Network Messages* originated from another process, or *Intra-process* events used by protocols to interact among each other within the same process.

Each protocol is exclusively assigned a dedicated thread, which handles received events in a serial fashion by executing their respective callbacks. In a single Babel process, any number of protocols may be executing simultaneously, allowing multiple protocols to cooperate (i.e., multi-threaded execution), while shielding developers from concurrency issues, as all communication between protocols is done via message passing.

From the developer's point-of-view, a protocol is responsible for defining the callbacks used to process the different types of events in its queue. The developer registers the callback for each type of event and implements its logic, while Babel handles the events by invoking their appropriate callbacks. While relatively simple, the event-oriented model provided by Babel allows the implementation of complex distributed protocols by allowing the developer to focus almost exclusively on the actual logic of the protocol, with minimal effort on setting up all the additional operational aspects.

2) *Core*: The Babel core is the central component which coordinates the execution of all protocols within the scope of a process.

As illustrated in Figure 1, every interaction in Babel is

mediated by the core component, as it is this component's responsibility to deliver events to each protocol's event queue. Whenever a protocol needs to communicate with another protocol, it is the core that processes and delivers events exchanged between them. When a message is directed to a protocol in another process, the core component delivers it to the network channel used by the protocol, which then sends the message to the target network address. That message is then handled by the core on the receiving process that ensures its delivery to the correct protocol.

Besides mediating interaction between protocols (both inter and intra process), the core also keeps track of timers setup by protocols, and delivers an event to a protocol whenever a timer setup by the protocol is triggered.

3) *Network*: Babel employs an abstraction for networking which we name *channels*. Channels abstract all the complexity of dealing with networking, and each one provides different behaviors and guarantees. Protocols interact with channels using simple primitives (`openConnection`, `sendMessage`, `closeConnection`), and receive events from channels whenever something relevant happens. These events are channel-specific and are handled by protocols just like any other type of event (i.e., by registering a callback for each relevant channel event).

For instance, the framework provides a simple `TCPChannel` which allows protocols to establish and accept TCP connections to/from other processes. This channel generates events whenever an outgoing connection is established, fails to be established, or is disconnected, and also whenever an incoming connection is established or disconnected. Other examples of provided channels include a channel with explicit and automatic acknowledgement messages, a channel that creates one connection for each protocol running in different processes, and a `ServerChannel` that does not establish connections, only accepts them, and its corresponding counter-part, the `ClientChannel`. We also provide a TCP-based channel that implements the φ -accrual failure detector [25], which notifies protocols that registered a callback whenever another process is suspected.

The framework also allows developers to design their own channels if they need to enforce some specific behavior or guarantee at the network level for a protocol to function correctly. Network channels are implemented using Netty [26], which is a popular Java networking framework. However, the typical developer of Babel does not have to interact with Netty directly.

A protocol can use any number of channels, and a channel can be shared by more than one protocol. In the example of Figure 1, two channels were instantiated by Babel. Channels within a Babel process are instantiated on demand by the Babel core when protocols are instantiated. Upon protocol instantiation, protocols define the channels they will be using, instructing the Babel core to prepare and instantiate the necessary network channels.

A. API

Babel is provided as a Java library. Protocols in Babel are developed by extending an abstract class - `GenericProtocol`. This class contains all the required methods to generate events and register the callbacks to process received events. Each protocol is identified by a unique identifier, used to allow other protocols to interact with it. There is also a special *init* event that protocols must implement, which is usually employed to define a starting point for the operation of the protocol (e.g., communicate with some contact node already in the system or setup a timer event).

The API can be divided in three categories: timers, inter-protocol communication (within the same process), and networking. Tables I and II list, respectively, the API for active operations (i.e. operations that generate events) and reactive operations (i.e. operations that define how to handle events).

1) *Timers*: Timers are essential to capture common behaviors of distributed protocols. They allow the execution of periodic actions (e.g., periodically exchange information with a peer), or to conduct some action a single time in the future (e.g., define a timeout).

In Babel, using timers can be achieved as follows. First, the developer needs to create a Java class that represents a timer, with a unique type identifier and extending the generic `ProtoTimer` class. Additionally, a timer might have any number of fields or logic as the developer needs. Listing 1 shows an example of the usage of timers in a protocol, with a timer that contains a counter that can be decremented (lines 1-11). To use the timer in a protocol, a callback method must be defined to be executed once the timer expires. This method must receive as parameters the timer object and its instance id, which is generated upon setting up the timer (lines 21-23). This callback is registered by calling the method `registerTimerHandler`, which takes as arguments the unique type identifier of the timer, and the callback itself (line 18). After registering the handler, any number of single-time or periodic timers can be setup using the methods `setupTimer` or `setupPeriodicTimer`, respectively. These methods take as parameters an instance of the timer object, and the delay to trigger the timer (in milliseconds). The periodic timer also receives a third parameter: the periodicity after the first timeout (line 19). Cancelling timers is also possible – for this, we simply call the method `cancelTimer` with the identifier of a previously setup timer as parameter. This identifier is obtained from the return value of the methods used to setup the timer or from the second parameter of the callback (line 22).

Listing 1: Timer Example

```

1 public class CounterTimer extends ProtoTimer {
2     public static final short TIMER_ID = 101;
3     int counter;
4
5     public CounterTimer(int initialValue) {
6         super(TIMER_ID);
7         counter = initialValue;
8     }
9
10    public int decCounter() {
11        return --counter;
12    }
13 }
14
15 public class CountdownProtocol extends GenericProtocol {

```

TABLE I: Babel Protocol Active API

	API	Description
<i>Timers</i>	setupTimer setupPeriodicTimer cancelTimer	Prepares a Timer event to be triggered after the defined time. Prepares a Timer event to be triggered periodically with a defined interval. Cancels a Timer event.
<i>Inter-protocol</i>	sendRequest sendReply triggerNotification	Sends a Request event to a specified protocol. Sends a Reply event to a specified protocol. Triggers a Notification event.
<i>Network</i>	createChannel openConnection closeConnection sendMessage	Creates a network channel. Opens a connection on a network channel to the given destination. Closes a connection on a network channel to the given destination. Send a Message event on a network channel to the given destination.

TABLE II: Babel protocol reactive API

	API	Description
<i>Timers</i>	registerTimerHandler	Registers the callback to use when a specific Timer event is triggered.
<i>Inter-protocol</i>	registerRequestHandler registerReplyHandle subscribeNotification unsubscribeNotification	Registers the callback to handle a given Request event. Registers the callback to handle a given Reply event. Subscribes to a given Notification event and registers the callback to handle it. Cancels a Notification subscription.
<i>Network</i>	registerMessageHandler registerMessageSerializer registerChannelEventHandler	Registers the callback to handle receiving a given Message event. Defines the serializer to use when encoding and decoding a given Message event. Registers the callback to use when a receiving a given channel-specific event.

```

13 public CountdownProtocol() {
14     super("CountdownProtocol", 100);
15 }
16
17 @Override
18 public void init(Properties props) throws HandlerRegistrationException {
19     registerTimerHandler(CounterTimer.TIMER_ID, this::handleCounterTimer);
20     setupPeriodicTimer(new CounterTimer(10, 1000, 300));
21 }
22
23 private void handleCounterTimer(CounterTimer timer, long timerId) {
24     if(timer.decCounter() == 0) cancelTimer(timerId);
25 }

```

2) *Inter-protocol communication*: Our framework supports multiple protocols executing concurrently in the same process. As such, we offer mechanisms for these protocols to interact with each other, allowing them to cooperate or delegate responsibilities. For instance, we could create a solution where a message dissemination protocol takes advantage of a peer-sampling protocol to obtain samples of the system’s filiation [9], [11], [27]. Listing 2 depicts such an example.

To support this, Babel provides two types of communication primitives: one-to-one requests/replies and one-to-many notifications. Similarly to timers, requests, replies, and notifications need to extend a generic class (ProtoReply, ProtoRequest, and ProtoNotification respectively) and have unique type identifiers. They can also have any extra arbitrary state and/or logic. Again, similarly to timers, we need to register a callback for each type of communication primitive (as seen in lines 7, 8, and 25). The callbacks all have the same parameters: the object that was sent and the identifier of the protocol who sent it. In order to send requests and replies, the methods `sendRequest` and `sendReply` are used. These methods take as parameters the Request or Reply to be sent and the destination protocol (as seen in lines 9 and 28 where, in the latter, the protocol replies to the sender of the request). For notifications however, the method `triggerNotification` (line 31) does not require

a destination, instead, every protocol that subscribed to that type of notification receives it (as exemplified in line 7).

Listing 2: Inter-protocol Communication Example

```

1 public class DisseminationProtocol extends GenericProtocol {
2     public DisseminationProtocol() {
3         super("DisseminationProto", 200);
4     }
5
6     @Override
7     public void init(Properties props) throws HandlerRegistrationException {
8         subscribeNotification(MembershipChangeNot.ID, this::onMembershipChange);
9         registerReplyHandler(MembershipReply.ID, this::onMembershipReply);
10        sendRequest(new MembershipRequest(), MembershipProtocol.PROTOCOL_ID);
11    }
12
13    private void onMembershipChange(MembershipChangeNot not, short emitter) {
14        updateMembership(not.getChanges());
15    }
16
17    private void onMembershipReply(MembershipReply reply, short from) {
18        setMembership(reply.getNodes());
19    }
20
21    public class MembershipProto extends GenericProtocol {
22        public static final short PROTOCOL_ID = 300;
23        public MembershipProto() {
24            super("MembershipProto", PROTOCOL_ID);
25        }
26
27        @Override
28        public void init(Properties props) throws HandlerRegistrationException {
29            registerRequestHandler(MembershipRequest.ID, this::onMembershipRequest);
30        }
31
32        private void onMembershipRequest(MembershipRequest request, short from) {
33            sendReply(new MembershipReply(currentMembership), sourceProto);
34        }
35
36        private void onConnectionEstablished(Peer peer) {
37            triggerNotification(new MembershipChangeNot(peer));
38        }
39    }

```

3) *Networking*: Naturally, as a framework for distributed protocols, Babel also provides abstractions to deal with networking (including management of connections). For this, we provide different network channels with different capabilities. The interaction of protocols with channels is (mostly) similar across different channels. Listing 3 presents an example of a protocol that sends a ping message to a random peer, waits

to receive a pong response message, and then chooses a new random peer to repeat the same behavior endlessly.

To use a channel, we start by setting up the properties for that channel. In the case of the TCPChannel, shown in the example, the required properties are the binding address and port for the listen socket (lines 9 to 11); other channels can consider different properties (e.g., a server channel that takes as property the maximum number of simultaneous client connections). Next, we create the channel by calling the method `createChannel`, passing the name of the channel and the properties object (line 12), and receiving an identifier representing the created channel. This identifier is useful if a protocol uses multiple channels simultaneously, to be able to select which channel to use to send a specific message, and to register different callbacks for different channels. Similarly to timers and requests/replies, we also need to create a class for each network message to be sent through a channel. Besides extending a generic class, and having a unique type identifier, the developer must also define a serializer for each message to enable the message to be encoded and decoded into network buffers. In lines 49-59 we show an example of a serializer for a message that only contains a single field. The serializer could be any function that converts Java objects to/from byte arrays, so one could use a JSON library or Java's own serializer (that however, would incur in additional overhead). We register the serializer for each message in a channel (since we could use different serializers for the same message across different channels) by using the method `registerMessageSerializer` (lines 13-14). We also need to register the callback for when a message is received, by using the method `registerMessageHandler` (lines 15-16). Note that different callbacks for the same message across different channels are supported. The callbacks to handle received messages take as arguments the message itself, the sender's IP address and port, the sender's protocol identifier, and the channel from which the message was received (as shown in lines 22 and 25).

Listing 3: Networking Example

```

1 public class PingPongProto extends GenericProtocol {
2     List<Host> peers;
3
4     public PingPongProto() {
5         super("PingPong", 400);
6     }
7
8     @Override
9     public void init(Properties props) throws HandlerRegistrationException {
10        peers = readPeersFromProperties(props);
11        Properties channProps = new Properties();
12        channProps.setProperty(TCPChannel.ADDRESS_KEY, props.getProperty("addr"));
13        channProps.setProperty(TCPChannel.PORT_KEY, props.getProperty("port"));
14        int cid = createChannel(TCPChannel.NAME, channProps);
15        registerMessageSerializer(cid, PingMsg.ID, PingMsg.serializer);
16        registerMessageHandler(cid, PongMsg.ID, PongMsg.serializer);
17        registerMessageHandler(cid, PingMsg.ID, this::uponPing);
18        registerChannelEventHandler(cid, PongMsg.ID, this::uponPong);
19        registerChannelEventHandler(cid, OutConnectionDown.EVENT_ID,
20            this::uponOutConnDown);
21        registerChannelEventHandler(cid, OutConnectionFailed.EVENT_ID,
22            this::uponOutConnFailed);
23        registerChannelEventHandler(cid, OutConnectionUp.EVENT_ID,
24            this::uponOutConnUp);
25        openConnection(chooseRandomPeer(peers));
26    }
27
28    private void uponPing(PingMsg msg, Host from, short sourceProto, int cid) {
29        sendMessage(new PongMsg(msg.getValue()), from, TCPChannel.CONNECTION_IN);
30    }
31
32    private void uponPong(PongMsg msg, Host from, short sourceProto, int cid) {
33        closeConnection(from);
34        openConnection(chooseRandomPeer(peers));
35    }
36 }

```

```

28 }
29 private void uponOutConnUp(OutConnectionUp ev, int cid) {
30     sendMessage(new PingMsg(System.currentTimeMillis()), ev.getPeer());
31 }
32 private void uponOutConnFailed(OutConnectionFailed<ProtoMessage> ev, int cid) {
33     openConnection(chooseRandomPeer(peers));
34 }
35 private void uponOutConnDown(OutConnectionDown ev, int channelId) {
36     openConnection(chooseRandomPeer(peers));
37 }
38 }
39 public class PingMsg extends ProtoMessage {
40     public final static short MSG_ID = 401;
41     private final long timestamp;
42
43     public PingMsg(long timestamp) {
44         super(MSG_ID);
45         this.timestamp = timestamp;
46     }
47
48     public int getTimestamp() {
49         return timestamp;
50     }
51
52     public static ISerializer<PingMsg> serializer = new ISerializer<>() {
53         @Override
54         public void serialize(PingMsg msg, ByteBuf out) throws IOException {
55             out.writeLong(msg.timestamp);
56         }
57         @Override
58         public PingMsg deserialize(ByteBuf in) throws IOException {
59             long timestamp = in.readLong();
60             return new PingMsg(timestamp);
61     };
62 }

```

A message can be sent using the `sendMessage` method, which takes as arguments the message to be sent, the destination address/port, the channel identifier (if more than one channel is being used) and optionally, the destination protocol. An additional parameter representing the connection to use can be passed. The interpretation of this parameter is however, channel dependent. Finally, each channel is responsible for generating notifications for relevant events that occur in it, for which the protocol can register callbacks. In the given example, the events generated by the TCPChannel (described before) are registered using the method `registerChannelEventHandler` (as seen in lines 17-19).

IV. EVALUATION

Our evaluation is based on two case studies of popular distributed applications. Our first case study is a simple peer-to-peer (P2P) application that disseminates messages to the network by leveraging two protocols, the HyParView [9] membership protocol that builds an unstructured overlay and provides nodes with stable neighboring nodes; and a Flood Dissemination protocol that disseminates a message to all neighboring nodes. Our second case study is a simple state machine replication (SMR) application, that leverages a MultiPaxos [18] consensus protocol to order the operations of an in-memory key-value store. The code for both case studies is available at <https://github.com/pedroAkos/babel-case-studies>.

Our evaluation is divided in three parts. In Section IV-A we provide implementation details over the protocols (implemented with Babel) for each case study. In Section IV-B, we perform a code review of publicly available MultiPaxos protocols and provide a comparison with our MultiPaxos implementation with Babel. Finally, in Section IV-C we present a performance evaluation on both case study applications.

A. Implementation Details

In this section we provide some implementation details over the protocols implemented with Babel for each case study. Each case study was implemented by a single PhD student familiarized with the protocols, and took about a work day to implement, analyze, and perform simple experiments to verify their operation and functionality.

a) P2P Application: The P2P application leverages two protocols (HyParView and a Flood Dissemination) that interact with each other through the inter-protocol communication mechanisms of Babel. In more detail, the Flood protocol registers interest in *neighbor up* and *neighbor down* notifications that are produced by the HyParView protocol whenever it connects to a new neighboring node, or disconnects from a neighboring node, respectively.

The HyParView protocol is implemented with 359 lines of Java code. This includes 15 event handlers, of which, 8 are network message handlers, one for each message processed by the protocol; 2 are timer handlers, one for a periodic action, and another for a timeout action; and 5 are channel event handlers, to the TCPChannel events. The events that support the operation of the protocol are implement with 361 lines of code, of which the majority (306) is dedicated to message events. Furthermore, a data structure implemented with 108 lines of code maintains the state of the protocol. In total, our implementation has 828 lines of code.

The Flood dissemination protocol is implemented with 74 lines of code. This includes 5 event handlers, of which 1 is a message handler that processes the broadcasted messages; 1 is a request handler that processes the request of the application to broadcast a message; and 2 are notification handlers that process the notifications triggered by HyParView. The protocol has 3 events that support its operation (a BroadcastMessage, a BroadcastRequest, and a DeliveryReply) that are implemented with 91 lines of code. We have also implemented a utility class that produces numeric hashes of the contents of the broadcasted messages to serve as unique message identifiers.

b) SMR Application: The SMR application is an in-memory key-value store that submits (client issued) operations to a consensus protocol for ordering. The consensus protocol exposes a notification that notifies the application of the execution order of operations. We have implemented two variants of MultiPaxos: a classic MultiPaxos variant where the acceptors inform all learners of their accepted values, and; a distinguished learner MultiPaxos variant, where the acceptors inform the leader (i.e., current proposer) of the accepted value, who then informs all learners.

The classic MultiPaxos is implemented with 235 lines of code, this includes the implementation of the proposer, acceptor, and learner components of the protocol. The protocol has 13 event handlers, where 4 are network message handlers, one for each protocol message; 3 are timer handlers for handling different timeout operations of the protocol (e.g., suspect the leader is dead, try to reconnect to a peer, enforce leadership); 1 is a request handler that handles the interface with the application(s); and finally, 5 are channel event handlers for

TABLE III: MultiPaxos Implementations Lines of Code

Implementation	Component	Java Lines of Code
Babel-MultiPaxos-Classic	Total	735
	Main Protocol Logic	235
	Events	250
	Utils	250
Babel-MultiPaxos-DistLearner	Total	787
	Main Protocol Logic	247
	Events	290
	Utils	250
MyPaxos	Total	1814
	Main Protocol Logic	683
	Messaging	306
	Utils	487
	Other	338
WPaxos	Total	22909
	Main Protocol Logic	2283
	Messaging	862 (+78 Protobuf)
	Networking	1376
	Support	536
	Utils	497
	Other	17891

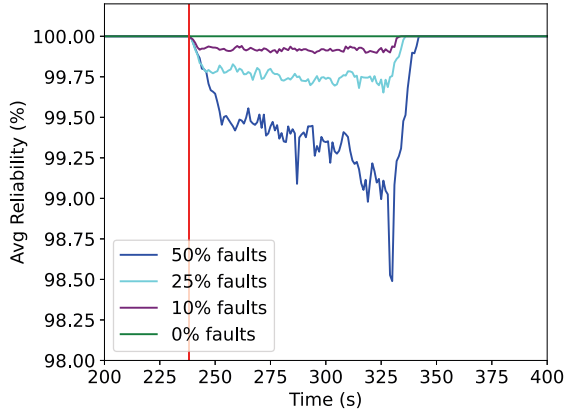
handling the TCPChannel events. The events of the protocol are implemented with 250 lines of code, where 163 lines of code are for messages. Furthermore, we have implemented some utility classes that help maintain the state of each consensus instance and operation, which were implemented with 250 lines of code. In total, our implementation has 735 lines of code.

The distinguished learner MultiPaxos is implemented with 247 lines of code. This implementation is very similar to the previous, with one additional network message handler for the message used to inform the learners of the accepted value.

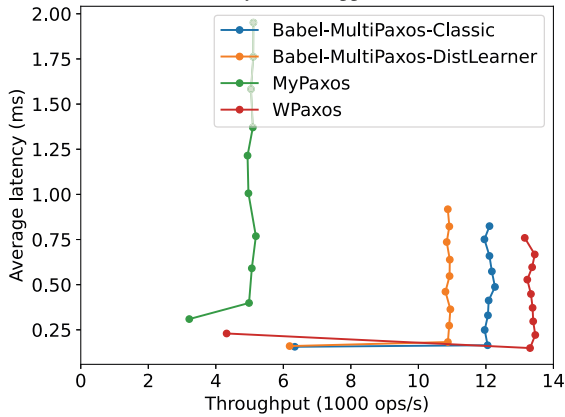
B. Code Review

In this section we provide a code review of publicly available MultiPaxos Java implementations, along with a comparison with our Babel MultiPaxos implementations. We chose the two MultiPaxos Java implementations on GitHub with the highest number of stars to perform this evaluation. The implementations are *MyPaxos* [28] and *WPaxos* [29], which had 131 and 130 stars at the time, respectively. In fact, these two implementations provide an interesting comparison as they, while implementing very similar MultiPaxos variants, have key differences that highlight some of the benefits of Babel. *MyPaxos* presents a very simple but naive implementation, while *WPaxos* presents a more complex and performant implementation. Despite this, they share the following design choices: *i*) the roles of proposer, acceptor, and learner are logically split across different Java classes; *ii*) instances are not executed concurrently (i.e., an instance only starts after the previous one is decided); and *iii*) every node that receives client operations competes to be the leader, meaning the application needs to redirect all operations to a single node to avoid poor performance and/or failed operations.

The only considerable design difference between the two implementations is that *MyPaxos* uses the classic message flow of MultiPaxos, while *WPaxos* uses the distinguished learner variant. In the following, we detail implementation aspects



(a) Reliability of P2P application.



(b) Throughput of Multipaxos implementations.

Fig. 2: Performance evaluation results

that were found while reviewing the code of each solution. Furthermore, Table III summarizes the number of lines of code for each implementation, where it is noticeable that our implementations have significant less lines of code.

a) MyPaxos: This implementation sacrifices performance and optimizations in favor of a small, simple-to-understand, and clear code structure. In short, in this implementation each node has three main threads, each corresponding to one of the Paxos roles. These communicate by sending messages (both to other nodes and to other roles within the same process). The implementation uses a simple networking library and relies on Java’s native object serialization to encode messages before sending them.

One thing we noted in this implementation, is that learners would periodically request learned values from other learners, to know when an instance had been decided. However, this heavily limited the rate at which operations were executed. As such, we made a small modification to the protocol, making learners eagerly broadcast their values as soon as they were accepted. This modification was done with about 10 lines of code without changing the implemented logic by the original

authors.

Additionally, we observed that this implementation was not very robust when tested under heavy load. The first issue we faced was a concurrency issue resulting in multiple threads accessing the same data structure, in this case a HashMap, which we solved by changing the data structure to the thread-safe ConcurrentHashMap. The second issue, resulted from a simplification made in the implementation that would spawn a new thread for each Paxos instance that would monitor a timeout. However, when executing the protocol under heavy load for a long period of time, the JVM crashed due to being unable to allocate memory for new threads. Babel aims at shielding the developer from issues such as these. In fact, Babel avoids concurrency issues by having sequential execution of events within a protocol, and its timer management allows protocols to register timers without the overhead of creating and managing extra threads.

In comparison with our Babel implementations, MyPaxos while being relatively simple, still requires more lines of code to implement than any of our Babel MultiPaxos implementations. This is mainly due to MyPaxos having to deal with the execution support code (e.g., threading, messaging, timers) that is handled by Babel.

b) WPaxos: This implementation is much more robust and performant than MyPaxos. It uses a custom communication layer, leveraging on the Netty [26] network framework, using TCP connections for large messages and UDP packets for small messages. Moreover, it serializes messages using Protocol Buffers [30]. While it also splits the Paxos roles in different Java classes, it handles the interactions between roles within the same process more efficiently than MyPaxos, by using direct method invocation instead of sending network messages. We were able to deploy and test this implementation without requiring any modification.

In comparison with our Babel implementations, it is clear that the use of Babel can greatly reduce the effort of programmers in creating working prototypes of distributed applications. While a considerable number of lines deal with features not implemented in our prototypes (mostly the 16275 lines of *Other* reported in Table III) such as persistency, there is still a substantial difference in the number of lines for the main protocol logic. This difference comes mostly from the logic dedicated to creating and handling timers, multiplexing received messages and triggered timers, as well as making sure received messages are handled sequentially. We note that using Babel all of these aspects do not have to be handled by the developer.

C. Performance Evaluation

In this section we describe our performance evaluation using the two case study applications. Our goal was, on the one hand to verify the correctness of our implementations and how they behave under faults (P2P application), and on the other hand to provide insights about the relative performance of these implementations when compared with standalone ones (SMR

application). In the following we describe the experimental setup and discuss the performance results for each case study.

a) P2P Application: Our experiments with the P2P application showcase a typical P2P evaluation with a network with a considerable number of nodes. In our experiments we measure the average reliability of message delivery in four different scenarios: *i)* a fault-free scenario; and fault scenarios where different fractions of nodes fail simultaneously *ii)* 10% faults; *iii)* 25% faults; and *iv)* 50% faults.

To execute our experiments we execute 100 docker containers distributed evenly across 4 identical servers with an AMD EPYC 7281 CPU and 128GiB of memory. Each container executes an instance of the application. We have applied latency among the containers with the Linux `tc` tool, based on a network generated with `inet` [31], having a mean latency of 293.39 milliseconds, to have a more realistic scenario. Each experiment runs for 12 minutes. Nodes start disseminating messages with a ratio of one message per second after two minutes since the start of the experiment until they have broadcasted a total of 500 messages. Faults are introduced at the middle of the experiment. Each experiment was executed three times. Results show the average of all runs for each experiment.

Figure 2a presents the results of all scenarios, with a vertical red line indicating the point of the experiments where faults were introduced. The x axis represents the time of the experiment when messages were being disseminated. The y axis represents the average reliability of message delivery. After the introduction of faults in the experiment, the average reliability drops marginally (up to 1%), this is because the overlay maintained by HyParView is very stable and maintains (high) connectivity. Nevertheless, there are a very few nodes that become temporarily disconnected and may not receive messages until they reconnect. The highest reliability drop in our experiment is when we fail 50% of the nodes in the network. This is to be expected as a lot of connections disappear. However, the HyParView protocol is resilient (reliability does not fall below 98%) to these faults as described by the authors.

b) SMR Application: Our experiments with the SMR application present an experimental comparison of different MultiPaxos implementations. In more detail, we compare our MultiPaxos implementations with the public ones described previously. To make comparisons fair, we adapted our key-value store application to work with MyPaxos and WPaxos. Furthermore, we disabled batching from all solutions, disabled the persistency in MyPaxos and configured WPaxos to perform the persistency asynchronously.

These experiments are executed in 4 identical servers with an Intel Xeon Gold 5220 CPU and 96GiB of memory, hosted in the Grid5000 [32] platform. Three servers execute the replicas of our application, while the last one executes the client. The client executes the YSCB [33] benchmarking tool, issuing operations in a closed loop to the key-value store application (which then orders them using the one of the studied MultiPaxos implementations) with a 50% read/write ratio. We vary the number of client threads from 1 to 10,

with increments of 1. Each experiment executes 5 times for a period of 90 seconds. We only consider data points from the 30 second mark to the 80 second mark to eliminate possible instabilities at the start or the end of an experiment. Results show the averages of each run for each implementation.

Figure 2b presents the average throughput (in the x axis) against the average latency (in the y axis) of operations for each implementation. We note that MyPaxos is the implementation with the least throughput. This is most likely due to the serialization used by the implementation which is known to have non-negligible overhead. In contrast, Babel MultiPaxos implementations show comparable performance results with the significantly more complex WPaxos implementation, which is able to achieve higher throughput due to the optimizations performed in the code albeit, at a much higher development overhead. We note that we validated the correctness of the solutions by inspecting the state of the replicated in-memory key-value store for each solution.

D. Discussion

Our evaluation validates three key aspects of Babel.

- The first is Babel’s generality as a tool to provide adequate support for implementing different classes of distributed protocols (either individually or in the form of a stack of protocols that interact with each others).
- The second key aspect relates to the correct/expected behavior of implemented protocols, which is validated by the first case-study, that displays results aligned with those published in the original paper.
- Finally, the third key aspect is the performance of implemented protocols. The performance of our Paxos implementation (second case-study), which was implemented with relatively low effort with Babel and has close performance to that of a highly optimized implementation of the protocol, show that there are no obvious and significant overhead sources from using the Babel framework.

V. CONCLUSION

In this paper we propose Babel, a Java framework to support the quick prototyping of distributed algorithms, particularly those that focus on fault-tolerance and dependability aspects of distributed systems. This tool can be used by researchers, practitioners, and teachers of distributed systems courses. Babel provides a set of abstractions that allows the developer to focus their efforts on implementing the algorithm logic, following an API that is close to the typical presentations either by descriptions or pseudo-code of such algorithms. Furthermore, Babel shields the developer from dealing with low-level networking aspects, as well as concurrency issues that arise when implementing protocols in a naive way, as shown in our comparison with public implementations of Paxos. We illustrated the benefits of Babel through two different case studies, and conducted a small experimental validation that shows that implementations that benefit from Babel present competitive performance when considering significantly more complex stand-alone implementations.

As future work, we plan to continue enriching Babel with novel network channels that can provide debugging tools that inspect the state of executing protocols and systems. We also plan on continuing to work on simpler and more powerful abstractions and interfaces to simplify the development of Babel protocols, along with integrated logging tools to facilitate debugging.

REFERENCES

- [1] N. A. Lynch, *Distributed Algorithms*. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 1996.
- [2] M. Raynal, *Fault-tolerant message-passing distributed systems: an algorithmic approach*. Springer, 2018.
- [3] C. Cachin, R. Guerraoui, and L. Rodrigues, *Introduction to Reliable and Secure Distributed Programming*, 2nd ed. Springer Publishing Company, Incorporated, 2011.
- [4] E. Rivière, “Simplifying hands-on teaching of distributed algorithms with splay,” in *2012 IEEE 26th International Parallel and Distributed Processing Symposium Workshops PhD Forum*, 2012, pp. 1311–1316.
- [5] L. Lamport, “The part-time parliament,” *ACM Trans. Comput. Syst.*, vol. 16, no. 2, p. 133–169, May 1998. [Online]. Available: <https://doi.org/10.1145/279227.279229>
- [6] L. Lamport et al., “Paxos made simple,” *ACM Sigact News*, vol. 32, no. 4, pp. 18–25, 2001.
- [7] R. Van Renesse and D. Altinbukan, “Paxos made moderately complex,” *ACM Comput. Surv.*, vol. 47, no. 3, Feb. 2015. [Online]. Available: <https://doi.org/10.1145/2673577>
- [8] A. Bessani, J. Sousa, and E. E. Alchieri, “State machine replication for the masses with bft-smart,” in *2014 44th Annual IEEE/IFIP International Conference on Dependable Systems and Networks*. IEEE, 2014, pp. 355–362.
- [9] J. Leitão, J. Pereira, and L. Rodrigues, “Hyparview: A membership protocol for reliable gossip-based broadcast,” in *Proc. of DSN’07*. IEEE, 2007.
- [10] S. Voulgaris, D. Gavidia, and M. Van Steen, “Cyclon: Inexpensive membership management for unstructured p2p overlays,” *Journal of Network and Systems Management*, vol. 13, no. 2, pp. 197–217, 2005.
- [11] M. Jelasity, S. Voulgaris, R. Guerraoui, A.-M. Kermarrec, and M. van Steen, “Gossip-based peer sampling,” *ACM Trans. Comput. Syst.*, vol. 25, no. 3, p. 8–es, Aug. 2007. [Online]. Available: <https://doi.org/10.1145/1275517.1275520>
- [12] Y. Kermarrec and L. Pautet, “Ada reusable software components for teaching distributed systems,” in *Software Engineering Education*, J. L. Díaz-Herrera, Ed. Berlin, Heidelberg: Springer Berlin Heidelberg, 1994, pp. 77–96.
- [13] A. Montresor and M. Jelasity, “Peersim: A scalable p2p simulator,” in *2009 IEEE Ninth International Conference on Peer-to-Peer Computing*. IEEE, 2009, pp. 99–100.
- [14] T. Issariyakul and E. Hossain, *Introduction to Network Simulator NS2*, 1st ed. Springer Publishing Company, Incorporated, 2010.
- [15] G. F. Riley and T. R. Henderson, “The ns-3 network simulator,” in *Modeling and Tools for Network Simulation*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2010, pp. 15–34.
- [16] M. Gupta, *Akka essentials*. Packt Publishing Ltd, 2012.
- [17] J. Armstrong, “erlang,” *Communications of the ACM*, vol. 53, no. 9, pp. 68–75, 2010.
- [18] L. Lamport et al., “Paxos made simple,” *ACM Sigact News*, vol. 32, no. 4, pp. 18–25, 2001.
- [19] H. Miranda, A. Pinto, and L. Rodrigues, “Appia, a flexible protocol kernel supporting multiple coordinated channels,” in *Proc. of ICDCS’01*, 4 2001, pp. 707–710.
- [20] S. Mena, X. Cuvellier, C. Gregoire, and A. Schiper, “Appia vs. cactus: comparing protocol composition frameworks,” in *Proc. of SRDS’03*, Oct 2003, pp. 189–198.
- [21] P. Á. Costa, A. Rosa, and J. Leitão, “Enabling wireless ad hoc edge systems with yggdrasil,” in *Proceedings of the 35th ACM/SIGAPP Symposium on Applied Computing (to appear)*, ser. SAC ’20. New York, NY, USA: ACM, 2020. [Online]. Available: <https://doi.org/10.1145/3341105.3373908>
- [22] W. Abdou, N. O. Abdallah, and M. Mosbah, “Visidia: A java framework for designing, simulating, and visualizing distributed algorithms,” in *2014 IEEE/ACM 18th International Symposium on Distributed Simulation and Real Time Applications*, 2014, pp. 43–46.
- [23] W. Schreiner, “A java toolkit for teaching distributed algorithms,” in *Proceedings of the 7th Annual Conference on Innovation and Technology in Computer Science Education*, ser. ITiCSE ’02. New York, NY, USA: Association for Computing Machinery, 2002, p. 111–115. [Online]. Available: <https://doi.org/10.1145/544414.544449>
- [24] P. Gouveia, J. a. Neves, C. Segarra, L. Liechti, S. Issa, V. Schiavoni, and M. Matos, “Kollaps: Decentralized and dynamic topology emulation,” in *Proceedings of the Fifteenth European Conference on Computer Systems*, ser. EuroSys ’20. New York, NY, USA: Association for Computing Machinery, 2020. [Online]. Available: <https://doi.org/10.1145/3342195.3387540>
- [25] N. Hayashibara, X. Defago, R. Yared, and T. Katayama, “The /spl phi/ accrual failure detector,” in *Proceedings of the 23rd IEEE International Symposium on Reliable Distributed Systems, 2004.*, 2004, pp. 66–78.
- [26] “Netty framework,” <https://netty.io/>.
- [27] J. Leitão, J. Pereira, and L. Rodrigues, “Epidemic broadcast trees,” in *Proc. of SRDS’07*. IEEE, 2007.
- [28] “Mypaxos,” <https://github.com/luohaha/MyPaxos>.
- [29] “Wpaxos,” <https://github.com/wuba/WPaxos>.
- [30] “Protocol buffers,” <https://developers.google.com/protocol-buffers>.
- [31] J. Winick and S. Jamin, “Inet-3.0: Internet topology generator,” Technical Report CSE-TR-456-02, University of Michigan, Tech. Rep., 2002.
- [32] “Wpaxos,” <https://www.grid5000.fr/>.
- [33] B. F. Cooper, A. Silberstein, E. Tam, R. Ramakrishnan, and R. Sears, “Benchmarking cloud serving systems with ycsb,” in *Proceedings of the 1st ACM symposium on Cloud computing*, 2010, pp. 143–154.