

Control Flow Versus Data Flow in Distributed Systems Integration: Revival of Flow-Based Programming for the Industrial Internet of Things

Wilhelm Hasselbring , Kiel University, 24118 Kiel, Germany

Maik Wojcieszak, wobe-systems GmbH, 24145 Kiel, Germany

Schahram Dustdar , TU Wien, 1040 Vienna, Austria

When we consider the application layer¹ of networked infrastructures, data and control flow are important concerns in distributed systems integration. Modularity is a fundamental principle in software design,² in particular for distributed system architectures. Modularity emphasizes high cohesion of individual modules and low coupling between modules. Microservices are a recent modularization approach with the specific requirements of independent deployability and, in particular, decentralized data management.³ Cohesiveness of microservices goes hand-in-hand with loose coupling, making the development, deployment, and evolution of microservice architectures flexible and scalable.⁴ However, in our experience with microservice architectures, interactions and flows among microservices are usually more complex than in traditional, monolithic enterprise systems, since services tend to be smaller and only have one responsibility, causing collaboration needs. We suggest that for loose coupling among microservices, explicit control-flow modeling and execution with central workflow engines should be avoided on the application integration level. On the level of integrating microservices, data-flow modeling should be dominant. Control-flow should be secondary and preferably delegated to the microservices. We discuss coupling in distributed systems integration and reflect the history of business process modeling with respect to data and control flow. To illustrate our recommendations, we present some results for flow-based programming in our Industrial DevOps project Titan, where we employ flow-based programming for the Industrial Internet of Things.

Data flow is concerned about where data are routed through a program/system and what transformations are applied during that journey. Control flow is concerned about the possible order of operations. These two concepts are somehow linked to each other: for example, the order of operations executed

in a computer program can influence where the data go. Similar, specific data values may steer the control flow. The side box below discusses data versus control flow in program analysis as an area related to this article.

DISTRIBUTED SYSTEMS INTEGRATION

File transfer, shared databases and Web resources, remote procedure calls, asynchronous messaging, and data streaming all support different forms of dataflow and control flow across distributed integrated systems

DATA VERSUS CONTROL FLOW PROGRAM ANALYSIS

In programming, when calling a function, starting the function's execution is control flow while passing the function's parameters is data flow. In this context, control and data flow are tightly linked, thus it is not straightforward to separate them:

- › Control-flow analysis deconstructs the order of operations in a computer program. This could be, for example, determining execution paths, but also precedence constraints between different operations.

The dominant question is how the locus of control moves through the program. Data may accompany the control flow, but is not dominant.

- › Data-flow analysis gathers information about the possible set of values calculated at various locations in a computer program.

The dominant question is how data moves through computations. As the data moves, control is activated.

Control flow refers to the path the execution takes in a program, and sequential programming that focuses on explicit control flow using control structures like loops or conditional statements is called **imperative programming**. In an imperative model, data may follow the control flow, but the main question is about the order of execution.

Dataflow abstracts over explicit control flow by placing the emphasis on the routing and transformation of data and is part of the **declarative programming** paradigm. In a dataflow model, control follows data and computations are executed implicitly based on data availability. Concurrency control refers to the use of explicit mechanisms like locks to synchronize interdependent concurrent computations. It is a matter of emphasis—control flow schedules data movement or data movement implies transfer of control.

with various degrees of coupling introduced in the integrated architecture.⁵

Service-oriented architecture (SOA) is an approach to developing enterprise systems by loosely coupling interoperable services from separate systems across different business domains. SOA emerged in the early 2000s, offering a way to develop new business services by reusing components from existing programs within the enterprise rather than writing functionally redundant code from scratch.

A crucial aspect of SOA is service orchestration. Developers utilize service orchestration to support the automation of business processes. Service orchestration is the coordination of multiple services exposed as a single aggregate service. In other words, service orchestration is the combination of service interactions to create higher level business services. This is usually accomplished through the use of a central workflow engine and/or an enterprise service bus (ESB). However, such a central orchestration service causes problems in microservice architectures.

- › A team should have full-stack responsibility for its microservices (Conway's Law). With a central orchestration service comes a central workflow team, which has to coordinate with the microservices teams whose services are involved in the workflow. With microservice architectures you create at

least one microservice per bounded context, according to domain-driven design.⁶ One important goal of microservices is to improve scalability and speed of the software development itself.⁴ Hence, it is common sense that one microservice needs to be owned by exactly one development team (which may own multiple microservices). Centrally managed ESBs do not fit into a microservices architecture. You may face a situation where you have to update your microservices in-sync with the central workflow model in case you make changes. This introduces a coupling between central and local control, which you do not want to have. However, inside a microservice or if a team owns multiple microservices, a workflow engine may be appropriate within this context.

- › For independent scalability, microservices should manage their data themselves; thus, manage their bounded context. Long running workflows need to keep persistent state somehow. This may imply coupling between the central orchestration service and the individual microservices. With microservices, workflows should only live inside service boundaries, if loose coupling is pursued.

Thus, we observe a conflict between central control (via orchestration) and independent evolution of microservices. This is particularly the case in our

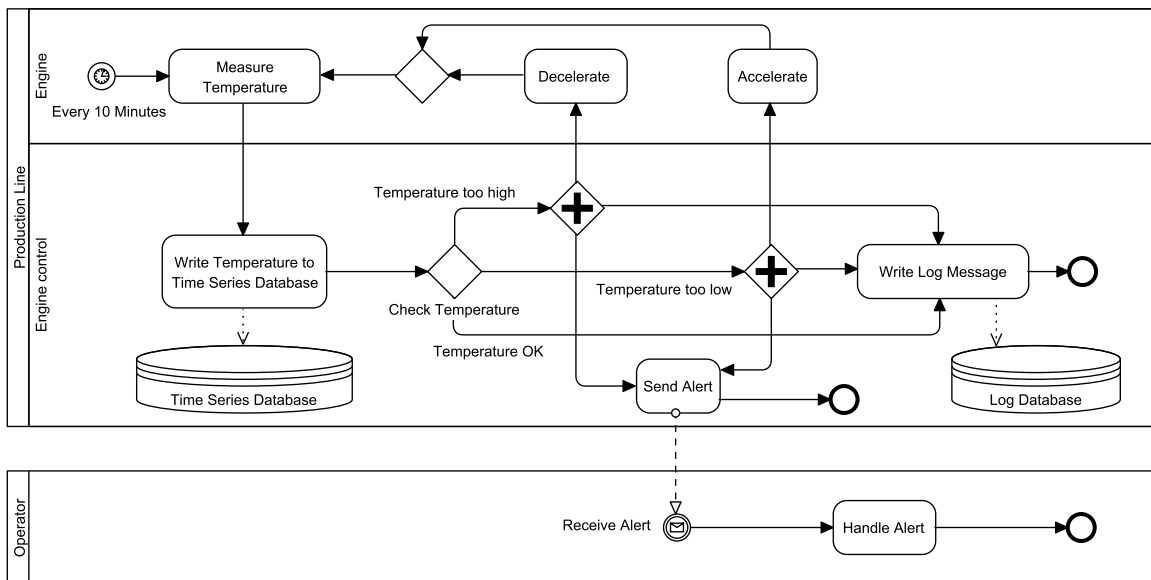


FIGURE 1. Example BPMN diagram for engine temperature control in a production line.

application domain of industrial automation, where for instance sensors and actuators are managed by third parties, with highly varying update cycles.

Service choreography is also related to service orchestration, as both are employed to create composite services and applications in SOAs; however, it is still worth pointing out the differences. A service choreography model works without a central orchestrator while a service orchestration model relies on a central controller to couple services. The microservice architectural style has promoted the idea of event-driven architectures to decouple your services. Smart endpoints and dumb pipes are preferred even more now than in previous generations of SOAs. This has not always been followed in all SOA implementations, resulting in ESB misuse.⁷

CONTROL FLOW MODELING

The workflow concept has de-facto become the standard paradigm for process modeling in business process management systems. Workflows are typically looked from three perspectives: 1) the control perspective, describing the logical order of tasks; 2) the data perspective, describing the information exchange between tasks; and 3) the resource perspective, describing the originators of tasks. Industry standards such as UML activity diagrams,⁸ the Business Process Model and Notation (BPMN)⁹ and event-driven process chains¹⁰ offer graphical notations for stepwise processes that include choice, iteration, and concurrent execution. However, data flow and control flow in

business workflows are not independent. The routing decisions in a workflow are typically based on data. The emphasis of these workflow modeling approaches is on control flow.

BPMN,⁹ for instance, does support data objects and data stores, so it is possible to use it to represent data flow, but control-flow modeling dominates BPMN models. Let us take a look at an illustrative example from our Industrial DevOps project.¹¹ Figure 1 shows an example BPMN workflow for temperature control of engines in a production line.

Timer events in BPMN are events which are triggered by a defined timer, in this example, a temperature sensor that periodically measures the engine's temperature. The engine control checks the temperature. If the temperature is too high or too low, the engine receives orders to decrease or increase the temperature, respectively. The measured temperatures are written to the time series database. Corresponding to the temperature check, appropriate messages are written to the log database. In case the temperature is too low or too high, the operator is alerted. The temperature is measured immediately after decelerating/accelerating the engine and periodically triggered by the timer event.

Check Temperature is a data-based exclusive gateway.⁹ BPMN offers several other gateways, including inclusive and parallel gateways. If, for instance, the temperature is too high, the production line receives the order to decrease the temperature, the corresponding log message is written, and the operator is alerted, all in parallel.

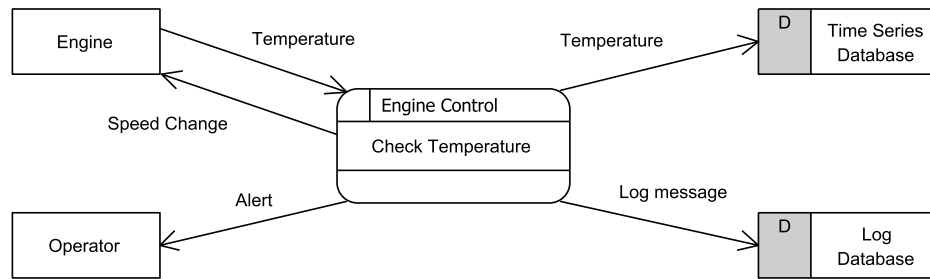


FIGURE 2. Example DFD for engine temperature control. Processes are depicted as squares with rounded corners, data store as open rectangles, external entities as closed rectangles, and data flows as directed lines.

The graphical, horizontal notation pools in Figure 1 depict the participants within the collaboration. Each pool forms a container for some processes. States in BPMN are linked by sequence, exception or message flows; sequence flows can be either incoming to or outgoing from a state. While *sequence* flows are restricted to an individual pool, *message* flows represent communications between pools. In Figure 1, alerting messages are exchanged between the Production Line and the Operator, depicted as a dashed arrow. In this example, no data objects are modeled, the emphasis is on control flow.

It is not apparent how to map parts of this model to bounded contexts and microservices. According to Evans,⁶ when a significant process or transformation in the domain is not a natural responsibility of an entity or value object, you should add an operation to the model as interface declared as a domain service. For instance, checking the temperature could be considered such a domain service. However, in Figure 1, this decision is modeled as control flow in the “central” workflow. As discussed in the previous section, we get a conflict between central control (via orchestration) and independent evolution of microservices when modeling this control flow on the integration level.

Our BPMN engine temperature control example is dominated by modeling the control flow. As an alternative approach, the JOpera Visual Composition Language¹² originates from the workflow area, but the JOpera approach emphasizes data flow. From the JOpera data flow graph, it is possible to derive the process’ control flow graph. JOpera includes a separate graph for modeling the control flow to specify control flow dependencies that cannot be automatically derived from the data flow. Before we take a look at such a combination of data and control flow, let us take a look at pure data flow modeling in the following section.

DATA FLOW MODELING

A data flow diagram (DFD) is a modeling technique for describing and analyzing information flows. It illustrates the flow of information based on input and output data. DFDs support structured analysis and design. They have the purpose of clarifying system requirements and major data transformations. DFDs illustrate business processes with the help of external data stores, the data flowing from one process to another, and delivering the result data. A DFD is a way to visualize the flow of data of a process or a system that aims to be accessible to both software engineers and domain experts alike. A DFD has no control flow, there are no decision rules and no loops.

Several DFD notations exist. We employ the Gane & Sarson notation.¹³ To facilitate the understanding of DFDs, the example DFD for engine temperature control in Figure 2 displays the four basic elements process, data store, data flow, and external entity, which are introduced as follows:

- ▶ Processes refer to the activities that operate the data of the system. A process receives input data and produces output with a different content or form. Processes can be as simple as collecting input data, or it can be complex as producing a report containing monthly sales. A process is depicted as squares with rounded corners with a unique name in form of verb or verb phrase, for example, Check Temperature in Figure 2. It is optional to indicate the place at which the process is executed, “Engine Control” in our example.
- ▶ Data stores represent the repository of data manipulated by processes, which can be databases or files (Time Series Database and Log Database in Figure 2). A data store is represented by an open rectangle in a DFD with a name in the form of noun or noun phrase. A data store is used

to represent a situation when the system must retain data because one or more processes need to use the stored data in a later time.

Note that such data stores could also be modeled with the BPMN. Our goal here is to explicitly illustrate the differences between control-flow modeling (with BPMN in Figure 1) and data-flow modeling (with DFD in Figure 2).

- › Data flows are directed lines indicating the data flow from or to a process with the information on the line of a data flow. At least one end of a data flow is linked to a process. Note that data cannot move without a process. In other words, data cannot go to or come from a data store or an external entity without having a process pushing it or pulling it. Data stores are passive while processes and external entities are active. In Figure 2, the engine sends the measured temperature to the Engine Control to Check the Temperature.
- › External entities are components outside of the boundaries of the modeled information system. They represent how the information system interacts with the outside world. Example external entities in Figure 2 are the Engine and the Operator. An external entity is depicted as a closed rectangle in DFDs. An external entity is a person, department, outside organization, or other information system that provides data to the system or receives outputs from the system.

Note that the DFD in Figure 2 does not describe exactly the same information as the control-flow BPMN diagram in Figure 1. Several (control-flow) details are left out, such as the logic to check the temperature. This decision is delegated to the appropriate domain service,⁶ for which a microservice will be responsible in a microservice architecture (this could be a domain service “Check Temperature” for our engine control example). Such control-flow concerns are delegated from the integration layer to the individual microservices, to reduce the coupling. Please note that just such a simple rule-based decision as in our illustrative example would be somewhat too fine grained to constitute a microservice.

TITAN FLOW-BASED PROGRAMMING

Titan is a software platform for integrating and monitoring industrial production environments, following

our Industrial DevOps approach.¹¹ With the Industrial DevOps approach, we intend to introduce methods and culture of DevOps into industrial production environments. The fundamental concept of this approach is a continuous process of development, operation, and observation of the entire production environment.

To achieve this, Titan applies the principles of flow-based programming. Flow-based programming is a programming paradigm, introduced by Morrison in the early 1970s.¹⁴ Flow-based programming defines applications as networks of black-box processes, which communicate via data traveling across predefined connections. As such, flow-based programming emphasize data flows, as the previously introduced DFDs do. Example Titan flows for engine temperature control are shown in Figure 3. In Titan, a graph of connected bricks is called a *flow*. There are several types of bricks¹⁵:

- › Inlets and outlets are to be found at the edges of flows and constitute the start and ending points of data flows (depicted as triangles embedded within squares). An inlet is a producer brick to a data flow. An outlet brick will end a data flow. In Figure 3, Engine is an inlet for temperature measurements. The Time Series Database Writer is an outlet to store the measurement data. Within Titan flows, communication among bricks is standardized, while inlets and outlets connect our flows to possibly heterogeneous external components.
- › Filter bricks (depicted as triangles, with the tip to the right) process the incoming data based on filter conditions. The Temperature Filter in Figure 3 converts the raw sensor measurements into corresponding Titan data structures.
- › Selector bricks (depicted as triangles, with the tip to the left) forward incoming data to selected out ports of this brick depending on the conditions set in the implementation. In Figure 3, the Check Temperature selector decides whether to handle the temperature or to write a log message.
- › Signals do constitute a special type of flow edge (depicted as double arrows to the left and right for signal producers and consumers, respectively). A signal producer starts a flow by receiving a trigger possibly including a dataset delivered via the signal. The signal consumer will end the current flow, optionally triggering some signal producers with the same name. In Figure 3, handle temperature is such a signal. With signals, we may build event-driven architectures.

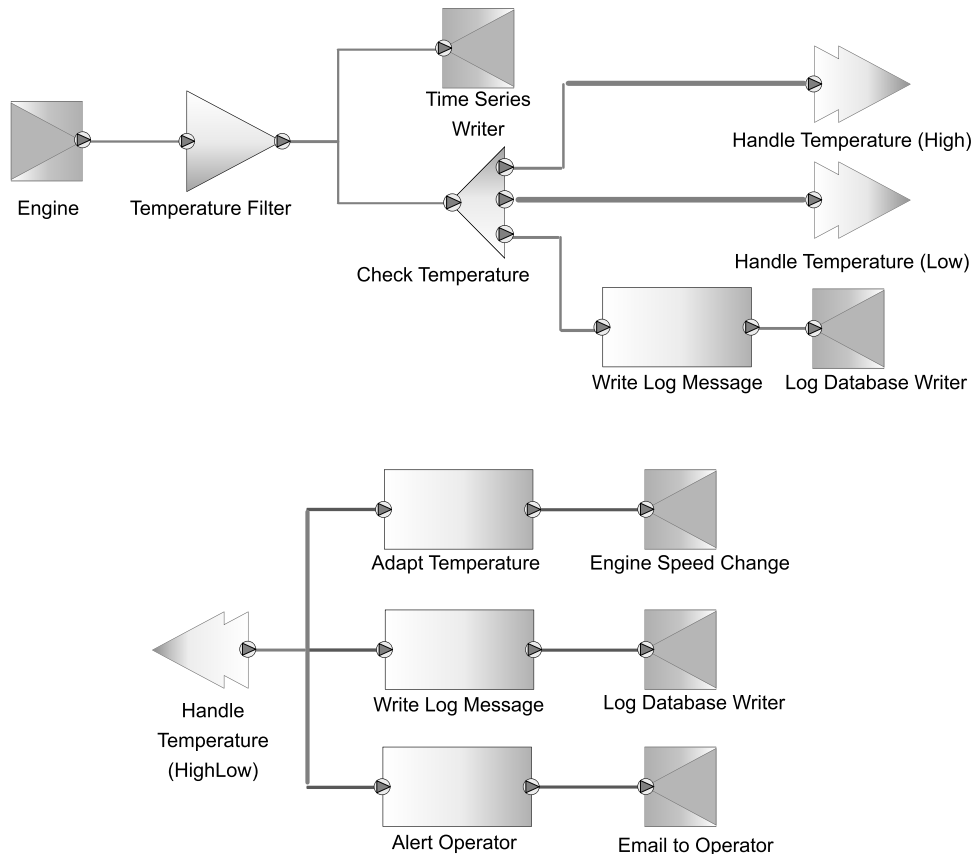


FIGURE 3. Example Titan flows for engine temperature control.

- ▶ General bricks (depicted as rectangles) contain logic that cannot be classified into the more specialized bricks listed above. In Figure 3, Adapt Temperature and Alert Operator are examples for such general bricks.

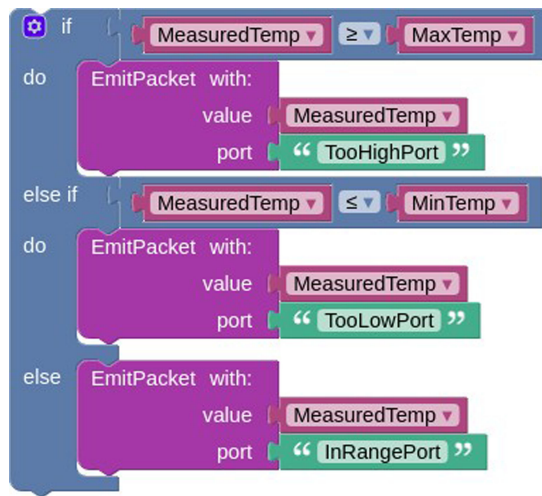
Note that the control flow decisions are encapsulated within the brick implementations. The visual Titan flows describe only data flow. The essential difference to the previous DFD model is that we model *event-driven architectures* with Titan’s signal mechanism. Event-driven architectures fit to microservice architectures.¹⁶

Titan provides a graphical modeling language, which is designed to enable the domain experts to model the integration and, based on this, to configure the integrated system. Hence, there is no software engineer required to perform these configurations or (preconfigured) changes. Instead, domain experts receive training on the modeling language, to allow for end-user programming.

The strong encapsulation of the internal functionality of a brick makes it highly modular. Bricks

implemented in different languages can be combined. The brick logic can be described in the form of a script, for example, with Python. With Titan, we also intend to combine our flow-based programming approach with block-based programming,¹⁷ such that the domain experts do not need to learn textual programming languages to describe the internal brick logic. In block-based programming, the programming constructs like conditionals and loops are represented via graphical blocks. Popular examples include MIT Scratch and Google Blockly. Figure 4 depicts a block-based program snippet and corresponding Python code for the Check Temperature selector brick of Figure 3.

With Titan, graphical flow-based programming for integrating distributed systems is dominated by data flows, while control flow is specified within the bricks. For both, we provide low-code programming to domain experts. In principle, the Titan approach is similar to JOpera,¹² where data flow graphs are refined by control flow graphs. However, with Titan, data flow graphs are refined via block-based programming to specify the control flow (see Figure 4).



```

if MeasuredTemp >= MaxTemp:
    EmitPacket(MeasuredTemp, 'TooHighPort')
elif MeasuredTemp <= MinTemp:
    EmitPacket(MeasuredTemp, 'TooLowPort')
else:
    EmitPacket(MeasuredTemp, 'InRangePort')

```

FIGURE 4. Block-based program snippet (left) and corresponding Python code (right) for the check temperature selector brick of Figure 3.

CONCLUSION

With the production line example, we intend to illustrate our experience and lessons learned with respect to control and data flow in distributed systems integration:

1. Our BPMN engine temperature control example in Figure 1 is a pure control-flow model. The logic for checking the temperature, for instance, is explicitly modeled in the (global) workflow.
2. Our DFD engine temperature control example in Figure 2 is a pure data-flow model. The logic for checking the temperature is *not* modeled in the (global) data flow, it should be implemented in a domain service.
3. Our Titan engine temperature control example in Figure 3 is a data-flow model, enriched with events (called signals in Titan). Such event-driven architectures fit well to microservice architectures and domain-driven design.¹⁶

Both data and control flow are important concerns in distributed systems integration. Based on our experience, we suggest that for loose coupling, explicit control-flow modeling should be avoided on the integration level. Modeling control flow is more coupled because it assumes an exact ordering of service invocations, while data flow abstracts from this ordering as long as the service interfaces have matching assumptions regarding the data that needs to be exchanged. Thus, on the level of integrating microservices, data-flow modeling should be dominant. Control-flow should be secondary and preferably delegated to the microservices in some way. However, be

aware of the resulting tradeoff of loosing an integrated overview on the control flow. Such an integrated overview on the actual system interactions may be reconstructed via runtime monitoring,^{18, 19} but the system design should focus on data flow. We suggest that researchers investigate more on data-flow oriented modeling methods, and that professionals reconsider and revive data-flow modeling and flow-based programming.

ACKNOWLEDGMENTS

This work was supported by the Federal Ministry of Education and Research (BMBF, Germany) in the Titan Project (<https://www.industrial-devops.org>) under Grant 01IS17084A/B.

REFERENCES

1. B. Zheng, J. Yin, S. Deng, Z. Wu, and S. Dustdar, "A service-oriented network infrastructure for crossover service ecosystems," *IEEE Internet Comput.*, vol. 24, no. 1, pp. 48–58, Jan. 2020.
2. D. L. Parnas, "On the criteria to be used in decomposing systems into modules," *Commun. ACM*, vol. 15, no. 12, pp. 1053–1058, Dec. 1972.
3. J. Lewis and M. Fowler, "Microservices," 2014. [Online]. Available: <https://martinfowler.com/articles/microservices.html>
4. W. Hasselbring and G. Steinacker, "Microservice architectures for scalability, agility and reliability in E-commerce," in *Proc. IEEE Int. Conf. Softw. Archit. Workshops*, Apr. 2017, pp. 243–246.

5. C. Pautasso and O. Zimmermann, "The web as a software connector: Integration resting on linked resources," *IEEE Softw.*, vol. 35, no. 1, pp. 93–98, Jan. 2018.
6. E. Evans, *Domain-Driven Design: Tackling Complexity in the Heart of Software*. Reading, MA, USA: Addison-Wesley, 2004.
7. D. Neri, J. Soldani, O. Zimmermann, and A. Brogi, "Design principles, architectural smells and refactorings for microservices: A multivocal review," *SICS Softw.-Intensive Cyber-Phys. Syst.*, vol. 35, no. 1/2, pp. 3–15, Aug. 2020.
8. M. Dumas and H. M. Arthur ter Hofstede, "UML activity diagrams as a workflow specification language," in *The Unified Modeling Language. Modeling Languages, Concepts, and Tools*, vol. 2185, pp. 76–90, New York, NY, USA: Springer, 2001, doi:10.1007/3-540-45441-1_7.
9. M. Chinosi and A. Trombetta, "BPMN: An introduction to the standard," *Comput. Standards Interfaces*, vol. 34, no. 1, pp. 124–134, Jan. 2012.
10. A.-W. Scheer, O. Thomas, and O. Adam, "Process modeling using event-driven process chains," in *Process-Aware Information Systems*, Hoboken, NJ, USA: Wiley, Oct. 2005, pp. 119–145, doi: 10.1002/0471741442.ch6.
11. W. Hasselbring, S. Henning, B. Latte, A. Möbius, T. Richter, S. Schalk, and M. Wojcieszak, "Industrial DevOps," in *Proc. IEEE Int. Conf. Softw. Archit. Companion*, Mar. 2019, pp. 123–126.
12. C. Pautasso and G. Alonso, "The JOpera visual composition language," *J. Vis. Lang. Comput.*, vol. 16, no. 1/2, pp. 119–152, Feb. 2005.
13. C. Gane and T. Sarson, *Structured Systems Analysis: Tools and Techniques*. Englewood Cliffs, N.J., USA: Prentice-Hall, 1979.
14. J. P. Morrison, *Flow-Based Programming: A New Approach to Application Development*. Scotts Valley, CA, USA: CreateSpace, 2nd ed. 2010. [Online]. Available: <http://www.jpaulmorrison.com/fbp/>.
15. T. Flow, 2020. [Online]. Available: <https://doc.industrial-devops.org/titanDocumentation/titan-flow-toc.html>.
16. E. Garofolo, "Practical Microservices: Build Event-Driven Architectures with Event Sourcing and CQRS," *The Pragmatic Programmers*, LLC, Sebastopol, 2020.
17. D. Bau, J. Gray, C. Kelleher, J. Sheldon, and F. Turbak, "Learnable programming: Blocks and beyond," *Commun. ACM*, vol. 60, no. 6, pp. 72–80, May 2017.
18. W. Hasselbring and A. van Hoorn, "Kieker: A monitoring framework for software engineering research," *Softw. Impacts*, vol. 5, Aug. 2020, Art. no. 100019.
19. W. Hasselbring, A. Krause, and C. Zirkelbach, "ExplorViz: Research on software visualization, comprehension and collaboration," *Softw. Impacts*, vol. 6, pp. 1–4, Nov. 2020, doi:10.1016/j.simpa.2020.100034.

WILHELM HASSELBRING is a Full Professor of Software Engineering with the Department of Computer Science, Kiel University, Kiel, Germany. Contact him at hasselbring@email.uni-kiel.de.

MAIK WOJCIESZAK is a Co-founder and CTO of wobe-systems GmbH, Kiel, Germany, providing software solutions for industrial automation. Contact him at mw@wobe-systems.com.

SCHAHRAM DUSTDAR is a Professor of Computer Science and the Head of the Distributed Systems Group, TU Wien, Vienna, Austria. He became Fellow of the IEEE in 2016 for contributions to elastic computing for cloud applications. He is the corresponding author of this article. Contact him at dustdar@dsg.tuwien.ac.at.