

Discovery of Functional Architectures From Event Logs

Jan Martijn E.M. van der Werf and Erwin Kaats

Department of Information and Computing Science
Utrecht University

P.O. Box 80.089, 3508 TB Utrecht, The Netherlands
j.m.e.m.vanderwerf@uu.nl, e.j.kaats@students.uu.nl

Abstract. The functional architecture focuses on decomposing functionality into modules that offer certain features. These features require interactions in order to complete their functionality. However, functional architectures typically only focus on the static aspects of the system design. Additional modeling techniques, such as like message sequence charts are often used in the early phases of software design to indicate how the software should behave.

In this paper we investigate the use of process discovery techniques to discover from these scenarios the internal behavior of individual components. Based on event logs, this paper presents an approach (1) to derive the information flows between features, (2) identify the internal behavior of features, and (3) to discover the order between features within a module. The approach results in a sound workflow model for each module. We illustrate the approach using a running example of a payment system.

1 Introduction

One of the principle tasks of a software architect is to design a software system [14], i.e., to organize the software elements the system is composed of in sets of structures, to allow reasoning about the system [4]. An important aspect of the software architecture is the functionality it offers. To decompose and specify the functionality of software, the authors of [6] introduced the Functional Architecture Model (FAM). It separates the functionality into features that are offered by the different modules the system is decomposed into. The informal notation of the model makes it a valuable discussion model between the people that define the functionality the software needs to offer, and the software architect [11].

Features interact with other features via information flows to offer their functionality. However, FAM only offers a static view on this interaction, i.e., the information flow only shows possible interactions, but imposes no order on or dependencies between these flows. Thus, to show how functionality is offered by the system, the architect requires additional models. One way is to define

scenarios on top of the models, in which the architect can specify which features interact in which order. These scenarios then result in event logs, that can be analyzed using process mining techniques [3]. Another source for discovering the possible interactions between features is the use of system execution data [], mapping events to the (partial) execution of features.

In this paper, we investigate the possibility to discover, given a set of scenarios in the form of an event log, the functional architecture of the system. We thereby focus on three basic questions on the functional architecture:

1. Which features interact?
2. What is the internal behavior of features?
3. What is the order in which features are executed within a module?

The first question focuses on the discovery of information flows: given an event log, is it possible to derive which features interact? As features use these information flows to offer their functionality, with the next question we investigate whether it is possible to derive the internal behavior of features, i.e., how does a feature use its information flows to complete its functionality? The last question deals with the high-level view of the functional architecture. To execute the system's functionality, the features within a module are called in a certain order. Can process discovery techniques be used to discover these orders?

The remainder of this paper is structured as follows. To illustrate the approach, Sec. 2 presents a running example which we will use throughout the paper. Next, Sec. 3 presents the basic notions used in the paper. Section 4 introduces the functional architecture model in more detail, after which in Sec. 5 we will focus on solving the three questions posed in the introduction. Section 6 concludes the paper.

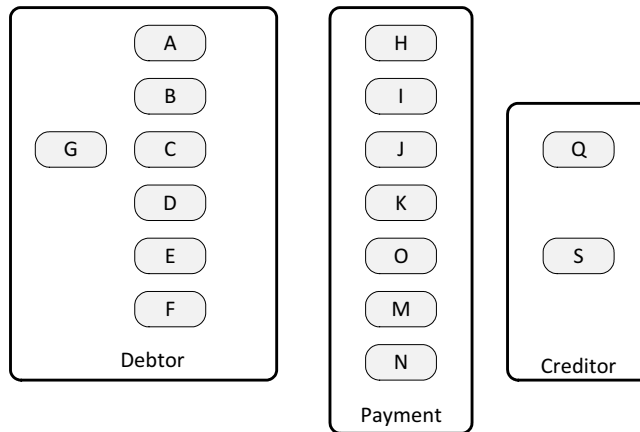


Fig. 1. Functional architecture model of the running example

2 Running Example

As an running example, consider the Payment System as introduced in [9]. The system consists of three modules, *Debtor*, *Payment* and *Creditor*. The payment module serves as an intermediate between the Debtor and the Creditor. An example of such a payment module is the european SEPA standard. The payment module initiates a transaction, which the debtor needs to accept. If the debtor accepts, the payment is continued, and the creditor is contacted to start the transaction. If for some reason the creditor rejects the transaction, the debtor is notified, and the transaction is terminated. Similarly, if the creditor accepts, the payment is passed to the debtor, and finally, the creditor receives the final payment information.

Unfortunately, no model exists that specifies the behavior of this system. Only, the system records the order in which the different features of the modules have been called in an event log, as shown in Tbl. 1. For readability, each of the features is abbreviated in this event log.

The *Debtor* module consists of the following features:

- Receive transaction request (A);
- Reject transaction (B);
- Accept transaction (C);
- Cancel transaction (D);
- Initiate payment (E);
- Send payment details (F);
- Archive transaction request (G).

Similarly, the *Payment* module consists of the following features:

- Send transaction request (H);
- Reject transaction request (I);
- Initiate creditor (J);
- Cancel transaction (K);
- Initiate payment (O);
- Handle payment (M);
- Archive transaction (N).

And the *Creditor* module consists of two features:

- Start transaction (Q);
- Handle transaction (S).

Based on this information, the architect can draw the modules with their features, as shown in Fig. 1. In the remainder of this paper, we investigate a method to use event logs, such as the one shown in Tbl. 1, to complete the diagram and derive a behavioral specification of the system.

Case	Trace
1	H, A, B, G, I, N
2	H, A, B, I, G, N
3	H, A, B, I, N, G
4	H, A, C, J, Q, S, K, D, G, N
5	H, A, C, J, Q, S, K, D, N, G
6	H, A, C, J, Q, S, K, N, D, G
7	H, A, C, J, Q, S, O, E, F, G, M, S, N
8	H, A, C, J, Q, S, O, E, F, G, M, N, S
9	H, A, C, J, Q, S, O, E, F, M, G, S, N
10	H, A, C, J, Q, S, O, E, F, M, G, N, S
11	H, A, C, J, Q, S, O, E, F, M, S, G, N
12	H, A, C, J, Q, S, O, E, F, M, S, N, G
13	H, A, C, J, Q, S, O, E, F, M, N, G, S
14	H, A, C, J, Q, S, O, E, F, M, N, S, G

Table 1. System execution data of the payment system

3 preliminaries

Let S be a set. The powerset of S is denoted by $\mathcal{P}(S) = \{S' \mid S' \subseteq S\}$. We use $|S|$ for the number of elements in S . Two sets U and V are *disjoint* if $U \cap V = \emptyset$. Some set S with relation \leq is a partial order, denoted by (S, \leq) , iff \leq is reflexive, i.e. $a \leq a$ for all $a \in S$, antisymmetric, i.e. $a \leq b$ and $b \leq a$ imply $a = b$ for all $a, b \in S$, and transitive, i.e. $a \leq b$ and $b \leq c$ imply $a \leq c$ for all $a, b, c \in S$. Given a relation $R \subseteq S \times S$ for some set S , we denote its transitive closure by R^+ , and the transitive and reflexive closure by R^* .

A *bag* m over S is a function $m : S \rightarrow \mathbb{N}$, where $\mathbb{N} = \{0, 1, \dots\}$ denotes the set of natural numbers. We denote e.g. the bag m with an element a occurring once, b occurring three times and c occurring twice by $m = [a, b^3, c^2]$. The set of all bags over S is denoted by \mathbb{N}^S . Sets can be seen as a special kind of bag where all elements occur only once; we interpret sets in this way whenever we use them in operations on bags. We use $+$ and $-$ for the sum and difference of two bags, and $=, <, >, \leq, \geq$ for the comparison of two bags, which are defined in a standard way.

A *sequence* over S of length $n \in \mathbb{N}$ is a function $\sigma : \{1, \dots, n\} \rightarrow S$. If $n > 0$ and $\sigma(i) = a_i$ for $i \in \{1, \dots, n\}$, we write $\sigma = \langle a_1, \dots, a_n \rangle$. The length of a sequence is denoted by $|\sigma|$. The sequence of length 0 is called the *empty sequence*, and is denoted by ϵ . The set of all finite sequences over S is denoted by S^* . We write $a \in \sigma$ if a $1 \leq i \leq |\sigma|$ exists such that $\sigma(i) = a$. *Concatenation* of two sequences $\nu, \gamma \in S^*$, denoted by $\sigma = \nu; \gamma$, is a sequence defined by $\sigma : \{1, \dots, |\nu| + |\gamma|\} \rightarrow S$, such that $\sigma(i) = \nu(i)$ for $1 \leq i \leq |\nu|$, and $\sigma(i) = \gamma(i - |\nu|)$ for $|\nu| + 1 \leq i \leq |\nu| + |\gamma|$. A sequence σ can be projected over some set U , denoted by $\sigma|_U$, and is inductively defined by $\epsilon|_U = \epsilon$, $(\langle a \rangle; \sigma)|_U = \langle a \rangle; \sigma|_U$ if $a \in U$, and $(\langle a \rangle; \sigma)|_U = \sigma|_U$ otherwise.

Petri Nets A *Petri net* [13] is a tuple $N = \langle P, T, F \rangle$ where (1) P and T are two disjoint sets of *places* and *transitions* respectively; and (2) $F \subseteq (P \times T) \cup (T \times P)$ is a *flow relation*. The elements from the set $P \cup T$ are called the *nodes* of N . Elements of F are called *arcs*. Places are depicted as circles, transitions as squares. For each element $(n_1, n_2) \in F$, an arc is drawn from n_1 to n_2 .

Let $N = \langle P, T, F \rangle$ be a Petri net. Given a node $n \in (P \cup T)$, we define its *preset* ${}^{\bullet}n = \{n' \mid (n', n) \in F\}$, and its *postset* $n^{\bullet} = \{n' \mid (n, n') \in F\}$. We lift the notation of preset and postset to sets. Given a set $U \subseteq (P \cup T)$, ${}^{\bullet}U = \bigcup_{n \in U} {}^{\bullet}n$ and $U^{\bullet} = \bigcup_{n \in U} n^{\bullet}$. If the context is clear, we omit the N in the subscript.

A *marking* of N is a bag $m \in \mathbb{N}^P$, where $m(p)$ denotes the number of *tokens* in place $p \in P$. If $m(p) > 0$, place p is called *marked* in marking m . A Petri net N with corresponding marking m is written as (N, m) and is called a *marked Petri net*. Given a marked Petri net (N, m) , transition t is enabled, denoted by $(N, m)[t]$, if ${}^{\bullet}t \leq m$. If transition t is enabled in (N, m) , it can fire, resulting in a new marking m' , denoted by $(N, m)[t](N, m')$, such that $m' + {}^{\bullet}t = m + t^{\bullet}$. We lift the firing of transitions to the firing of sequences in a standard way, i.e., a sequence $\sigma \in T^*$ of length n is enabled in (N, m) if markings m_0, \dots, m_n exist, such that $m = m_0$ and $(N, m_{i-1})[\sigma(i)](N, m_i)$ for all $1 \leq i \leq n$. A marking m' is reachable from some marking m in N , denoted by $(N, m)[*](N, m')$, if a firing sequence $\sigma \in T^*$ exists such that $(N, m)[\sigma](N, m')$. A marking m' is a *home marking* of (N, m) , if for all markings m'' with $(N, m)[*](N, m'')$, we have $(N, m'')[*](N, m')$.

A special class of Petri nets are the workflow nets [2]. A workflow net is a tuple $\langle P, T, F, i, f \rangle$ with $\langle P, T, F \rangle$ a Petri net, (2) $i \in P$ is the only place with no incoming transitions, (3) $f \in P$ is the only place with no outgoing transitions, i.e., ${}^{\bullet}i = f^{\bullet} = \emptyset$, and (4) all transitions have at least one incoming and one outgoing arc, i.e., ${}^{\bullet}t \neq \emptyset \neq t^{\bullet}$ for all $t \in T$.

Open Petri Nets Within a network of asynchronously communicating systems, messages are passed between the elements within the network. The approach we follow is based on Open Petri nets [5]. Communication in an open Petri net (OPN) is represented by special places, called the *interface places*. An interface place is either an *input place*, receiving messages from the outside, or an *output place* that sends messages to the outside of the OPN. An input place is a place that has only outgoing arcs, and an output place has no incoming arcs.

Definition 1. An Open Petri net is defined as an 7-tuple $\langle P, I, O, T, F, i, f \rangle$ where (1) $\langle P \cup I \cup O, T, F \rangle$ is a Petri net; (2) P is a set of internal places; (3) I is a set of input places, and ${}^{\bullet}I = \emptyset$; (4) O is a set of output places, and $O^{\bullet} = \emptyset$; (5) P, I and O are pairwise disjoint; (6) $i \in \mathbb{N}^P$ is the initial marking, and (7) $f \in \mathbb{N}^P$ is the final marking. We call the set $I \cup O$ the interface places of the OPN. An OPN is called closed if $I = O = \emptyset$.

An important behavioral property for OPNs is termination: an OPN should always have the possibility to terminate properly. We identify two termination properties: weak termination and soundness.

Definition 2. Let $\langle P, I, O, T, F, i, f \rangle$ be an OPN. It is weakly terminating, if f is a home marking of the marked Petri net $(\langle P \cup I \cup O, T, F \rangle, i)$. It is sound, if f is a home marking of the marked Petri net $(\langle P, T, F \rangle, i)$.

Communication between OPNs is done via the interface places. Two OPNs can only communicate if the input places of the one are the output places of the other, and vice versa.

Definition 3. Two OPNs A and B are composable, denoted by $A \oplus B$, if and only if $(I_A \cap O_B) \cup (O_B \cap I_A) = (P_A \cup T_A \cup I_A \cup O_A) \cap (P_B \cup T_B \cup I_B \cup O_B)$.

If A and B are composable, they can be composed into a new OPN, denoted by $A \oplus B$, with $A \oplus B = \langle P, I, O, T, F, i, f \rangle$ where $P = P_A \cup P_B \cup G$; $I = (I_A \cup I_B) \setminus G$; $O = (O_A \cup O_B) \setminus G$; $T = T_A \cup T_B$; $F = F_A \cup F_B$; $i = i_A + i_B$; and $f = f_A + f_B$ with $G = (I_A \cap O_B) \cup (O_B \cap I_A)$.


Event Logs and Behavioral Profiles Although event logs are defined as a tuple consisting of a set of case identifiers, events, and an attribute mapping [3], it is in this paper sufficient to consider an event log, denoted by \mathcal{L} , as a set of sequences over some alphabet T , i.e., $\mathcal{L} \subseteq T^*$. Given an event log \mathcal{L} , we define the *successor relation* [16] by $a <_{\mathcal{L}} b$ if a sequence $\sigma \in \mathcal{L}$ and $1 \leq i \leq |\sigma|$ exist, such that $\sigma(i) = a$ and $\sigma(i+1) = b$. Using the successor relation, we define the *behavioral profile* $(\rightarrow_c, \parallel_c, +_c)_{\mathcal{L}}$ as three relations: (1) the causality relation \rightarrow_c is defined by $a \rightarrow_c b$ iff $a <_{\mathcal{L}} b$ and $b \not<_{\mathcal{L}} a$, (2) the concurrency relation \parallel_c , which is defined by $a \parallel_c b$ iff both $a <_{\mathcal{L}} b$ and $b <_{\mathcal{L}} a$, and (3) the exclusive relation $+_c$ is defined by $a +_c b$ iff both $a \not<_{\mathcal{L}} b$ and $b \not<_{\mathcal{L}} a$ [16]. If the context is clear, we omit the subscript.

Given a marked Petri net (N, m) with $N = \langle P, T, F \rangle$, an event log $\mathcal{L} \subseteq T^*$ is called *complete* with respect to (N, m) iff traces $\sigma_1, \sigma_2 \in T^*$ exist such that $(N, m)[\sigma_1; \langle a, b \rangle; \sigma_2](N,)$ implies $a <_{\mathcal{L}} b$ for all $a, b \in T$.

4 Functional Architectures

To model the overview of a system, the modules it consists of, and the features these modules offer, we propose the use of the *functional architecture model* (FAM). The functional architecture of a system is “an architectural model which represents at a high level the software products major functions from a usage perspective, and specifies the interactions of functions, internally between each other and externally with other products” [6]. It offers modules containing features. Features of different modules interact via so-called *information flows*.

An example is shown in Fig. 2(a). The FAM contains 1 context module, E , 7 modules, A, B, C, D, X, Y and Z . Modules have features, depicted by the rounded rectangles. For example, module C contains two features, K and L . Between features of different modules, information flows exist, e.g., the information flow (F, q, L) between modules A and C .



(a) Example Functional Architecture Model

(b) Scenario as an overlay

Definition 4. A Functional Architecture Model (FAM) is defined as a 6-tuple $\langle \mathcal{M}, \mathcal{F}, h, m, \rightarrow \rangle$ where

- \mathcal{M} is a finite set of modules;
- \mathcal{C} is a finite set of context modules;
- \mathcal{F} is a finite set of features;
- $h : \mathcal{M} \rightarrow \mathcal{M}$ is the hierarchy function, such that the transitive closure h^* is irreflexive;
- $m : \mathcal{F} \rightarrow \mathcal{M} \cup \mathcal{C}$ is a feature map that maps each feature to a module, possibly in the context, and this module does not have any children, i.e. $h^{-1}(m(F)) = \emptyset$ for all $F \in \mathcal{F}$;
- $\rightarrow \subseteq \mathcal{F} \times \Lambda \times \mathcal{F}$ is the information flow, such that for $(A, l, B) \in \rightarrow$ we have $m(A) \neq m(B)$. The labels for the information flows are unique per feature, i.e., (A, l, B) and (A, l, C) imply $B = C$ for all labels $l \in \Lambda$ and $(A, l, B), (A, l, C) \in \rightarrow$.

Although the information flows define the possible interactions between modules, it remains a static overview of the system. Therefore, one can use scenarios on top of the functional architecture, e.g. by creating an overlay, highlighting the information flows that are executed and the order in which they should occur. Formally, we represent a scenario as a partial order.

Definition 5. Let $F = \langle \mathcal{M}, \mathcal{F}, h, m, \rightarrow \rangle$ be a FAM. A scenario of F is a pair $(S, <)$ with $S \subseteq \rightarrow$, such that (S, \leq) with $\leq = <^*$ is a partial order.

An example is shown in Fig. 2(b). The scenario implied by the overlay can be represented by a partial order induced by $(O, p, F) < (F, q, L)$, $(F, q, L) < (K, s, H)$, $(F, q, L) < (K, r, N)$, $(K, r, N) < (N, u, H)$, $(K, s, H) < (H, t, G)$, $(N, u, H) < (H, t, G)$, $(K, r, N) < (M, v, H)$, and $(M, v, H) < (H, t, G)$.

However, such scenarios are typically not specified. Another important drawback of such scenarios is their analyzability. Although each scenario can be

checked, the consistency between the different scenarios remains a difficult task. Therefore, in the remainder of this paper, we search for a method to derive the behavioral specification as a network of asynchronously communicating systems, given the system execution data produced by the actual system in the form of event logs.

5 Discovery of a Functional Architecture

In this section, we study the possibilities process mining [3] offers to generate Petri nets for each of the different modules a system consists of. Event logs describe the order in which features of a system have been executed. Such event logs are system wide. Instead of each module having its own event log, only global sequences exist, i.e., sequences concatenate the executed features over all modules. We assume that each feature belongs to exactly one module, and that the direction of communication between two features is in one direction, i.e., given two features A and B , we assume that either A sends a message to B , or vice versa, that B sends a message to A , but not both.

The behavioral specification of a system is three-fold: (1) communication between modules via their features, (2) the internal behavior within each feature, and (3) the order in which features are called within a module. In this section, we explore all three types of behavioral specification to come to a composed system of asynchronously communicating systems.

In the remainder, let \mathcal{L} be an event log over a set of features T , and let $\mathfrak{R} : T \rightarrow M$, with M the set of modules, be a function that maps each feature onto the module that contains that feature.

5.1 Communication between Features

Communication between modules within a system is asynchronous of nature: messages are sent between features in order to complete their functionality. Within an event log, we need to consider the order in which events or features occur. For example, given some trace σ , if the resource is different for two subsequent events, i.e., $\mathfrak{R}(\sigma(i)) \neq \mathfrak{R}(\sigma(i+1))$, then this might indicate that the former sends a message to the latter. This is expressed by the communication successor.

Definition 6 (Communication successor). *Let $\mathcal{L} \subseteq T^*$ be an event log. We define the communication successor relation $\ll_{\mathcal{L}} \subseteq T \times T$ by $A \ll_{\mathcal{L}} B$ iff $\mathfrak{R}(A) \neq \mathfrak{R}(B)$, $\sigma(i) = A$, and $\sigma(i+1) = B$ for some $\sigma \in \mathcal{L}$ and $1 \leq i < |\sigma|$.*

Although at first sight the communication successors seem to work, we need to remember the concurrent nature of asynchronous communication. Consider for example the communication between modules M and N as depicted in Fig. 2. We have $A \ll E$, which is indeed the communication as modeled in the composition $M \oplus N$. However, we also find $G \ll B$, indicating a possible communication between G and B . Listing all communication successors, we get $A \ll E$, $G \ll B$,




Fig. 2. Modules M and N

Case	Trace
1	A, E, F, G, B, C
2	A, E, F, B, G, C
3	A, E, B, F, G, C
4	A, E, B, F, C, G
5	A, E, F, B, C, G
6	A, B, E, F, G, C
7	A, B, E, F, C, G

Table 2. Corresponding event log

	E	F	G
A	\rightarrow	+	+
B			
C	+	\leftarrow	

Table 3. Communication behavioral profile

$F \ll B$, $B \ll G$, $G \ll C$, $E \ll B$, $B \ll F$, $F \ll C$, $C \ll G$, and $B \ll E$. Observe that because of the asynchronous nature of the communication, features B and E are concurrently enabled in Fig. 2. If the event log is complete, as is the case in this example, this should become visible in the communication successor relation, as for the normal successor relation on event logs. Thus, calculating the behavioral profile of the communicating transitions using the communication successor relation, we obtain the relations as shown in Tbl. 3. Indeed, the behavioral profile shows that B and E are concurrently enabled. Following the behavioral profile, we see that the causal relation of the behavioral profile correctly identifies the feature communication.

Using the communication behavioral profile, we can construct the information flows from an event log as follows. If $A \rightarrow B$ in the communication behavioral profile of the event log, then an information flow (A, x, B) exists, with x a fresh label. This results in the following translation:

Definition 7 (Generated FAM). Let \mathcal{L} be an event log, and $\{\rightarrow_c, ||_c, +_c\}$ be its communication behavioral profile. Its corresponding functional architecture model $\langle \mathcal{M}, \mathcal{F}, h, m, \rightarrow \rangle$ is defined by:

- $\mathcal{M} = \mathfrak{R}(\mathcal{L})$;
- $\mathcal{C} = \emptyset$;
- $\mathcal{F} = T$;
- $h = \emptyset$;
- $m = \mathfrak{R}$; and
- $\rightarrow = \{(A, x, B) \mid A \rightarrow_c B, \text{ and } x \in \Lambda \text{ a fresh label}\}$.

	Debtor							Payment							Creditor	
	A	B	C	D	E	F	G	H	I	J	K	O	M	N	Q	S
A								←	+	+	+	+	+	+	+	+
B								+	→	+	+	+	+	+	+	+
C								+	+	→	+	+	+	+	+	+
D								+	+	+	←	+	+		+	+
E								+	+	+	+	←	+	+	+	+
F								+	+	+	+	+	→	+	+	+
G								+		+	+	+			+	
H	→	+	+	+	+	+	+								+	+
I	+	←	+	+	+	+									+	+
J	+	+	←	+	+	+	+								→	+
K	+	+	+	→	+	+	+								+	←
O	+	+	+	+	→	+	+								+	←
M	+	+	+	+	+	←									+	→
N	+	+	+		+	+									+	
Q	+	+	+	+	+	+	+	+	+	←	+	+	+	+		
S	+	+	+	+	+	+		+	+	+	→	→	←			

Table 4. Communication behavioral profile for the running example

After constructing the communication behavioral profile for the running example, shown in Tbl. 4, we can complete the functional architecture model. Based on the given system execution data, we see for example that feature H communicates with feature A , and feature S sends messages to features K and O , and receives messages from feature M . The complete functional architecture of the running example is shown in Fig. 3.

5.2 Internal Behavior of Features

As can be seen in the running example, features can send and receive multiple messages. For example, feature S sometimes sends a message to feature K and sometimes to feature O . Therefore, the next step in discovering the functional architecture is to reconstruct the internal behavior of each of the features. For this, we create for each of the features an event log, containing the features that it communicates with. We call this the corresponding *feature log*.

Definition 8 (Feature log). Let $\mathcal{L} \subseteq T^*$ be an event log, and let $F \in T$ be some feature. Let $(\rightarrow_c, \parallel_c, +_c)$ be the corresponding communication behavioral profile. The feature log \mathcal{L}_F is defined by $\mathcal{L}_F = \{\sigma|_{C(F)} \mid \sigma \in \mathcal{L}, F \in \sigma\}$ where $C(F) = \{A \mid A \rightarrow_c F \vee A \rightarrow_c F\}$.

Consider for example feature F in the running example. This feature communicates with features K , O and M , i.e., $C(F) = \{K, O, M\}$. Its feature log is the projection of the log on these features, i.e., $\mathcal{L}_F = \{\langle K \rangle, \langle O, M \rangle\}$.

On these feature logs, we apply the genetic tree miner, that always returns a sound workflow net (cf. [7, 10]). Next, we transform the discovered workflow

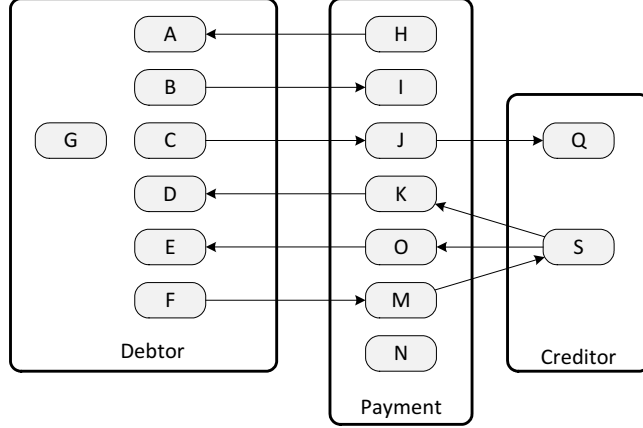


Fig. 3. Functional architecture model of the running example

net into an open Petri net, to visualize the messages sent and received by the feature. This results in a *feature net* for each of the features present in the event log.

Definition 9 (Feature Net). Let $\mathcal{L} \subseteq T^*$ be an event log, and let $F \in T$ be some feature. Let $(\rightarrow_c, \parallel_c, +_c)$ be the corresponding communication behavioral profile. The Feature net \mathcal{N}_F is the OPN $\langle P, I, O, T, F, i, f \rangle$ defined by

- $P = \bar{P}, T = \bar{T}, i = [\bar{i}], f = [\bar{f}]$;
- $I = \{p_{A-F} \mid A \rightarrow_c F\}$;
- $O = \{p_{F-A} \mid F \rightarrow_c A\}$;
- $F = \bar{F} \cup \{(t, p_{F-A}) \mid t \in T, \lambda(t) = A, F \rightarrow_c A\} \cup \{(p_{A-F}, t) \mid t \in T, \lambda(t) = A, A \rightarrow_c F\}$.

where $\langle \bar{P}, \bar{T}, \bar{F}, \bar{i}, \bar{f} \rangle$ is the discovered workflow net.

In our running example, each of the 16 features are transformed into a feature net. Most of the features are simple, like for feature H and A , consisting of a single transition sending a message to A , and receiving a message from H , respectively. A more complex feature net is the net for feature S , which internally decides whether it sends a message to K or to O . Figure 4 depicts some of the feature nets generated using the genetic tree miner.

5.3 Feature Interaction Within Modules

Now that each feature has its internal behavior defined by means of a feature net, the next step is to determine the order in which features are executed within each of the modules. As for the features, we first create event logs for each of the modules, by filtering each trace on the features it contains. This results in a *module log* for each of the modules.

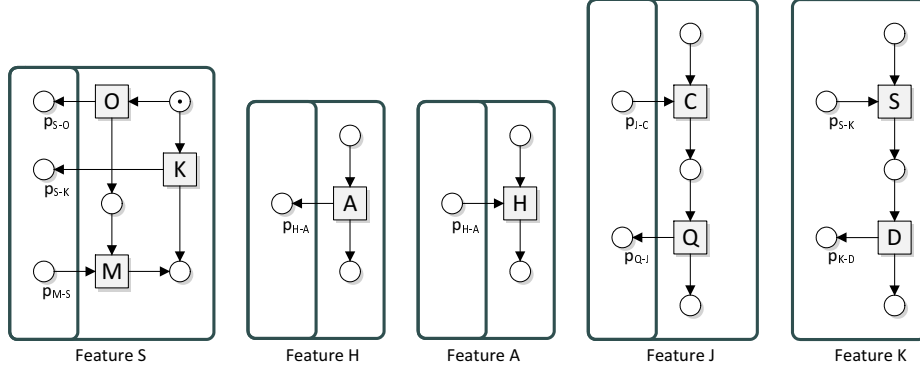


Fig. 4. Some of the feature nets generated for the running example

Definition 10 (Module Log). Let $\mathcal{L} \subseteq T^*$ be an event log. Let $M \in \text{Rng}(\mathfrak{R})$ be a module. Let $(\rightarrow_c, \parallel_c, +_c)$ be the corresponding communication behavioral profile. The Module log \mathcal{L}_M is defined by $\mathcal{L}_M = \{\sigma \mid \{F \mid \mathfrak{R}(F) = M\} \mid \sigma \in \mathcal{L}\}$.

Within the running example, we obtain three module logs, one for each of the modules. For example, module *Debtor*, has module log $\mathcal{L}_{\text{Debtor}} = \{\langle A, B, G \rangle, \langle A, C, D, G \rangle, \langle A, C, E, F, G \rangle\}$, and for *Creditor* we have $\mathcal{L}_{\text{Creditor}} = \{\langle Q \rangle, \langle Q, S, S \rangle\}$. Applying the inductive miner [10] results in the three workflow nets as depicted in Fig. 7. Notice that, although feature *S* occurs twice in one of the sequences, the algorithm only adds a single feature *S* in the resulting workflow model.

5.4 Composition of Feature Nets and Module Nets

Last step in the process is to combine the feature nets generated for each of the features with the generated module nets. This results in an open Petri net for each of the modules, defining the interaction between the different modules.

In the module net, each feature is represented by a single transition. Next step is to refine each feature by its feature net. For this, we first define the refinement of a transition by a workflow model on open Petri nets, as shown in Fig. 5. This refinement connect each input place of the refined transition with each of the transitions in the postset of the initial place of the workflow, and




Fig. 5. Refinement of a transition by a workflow net

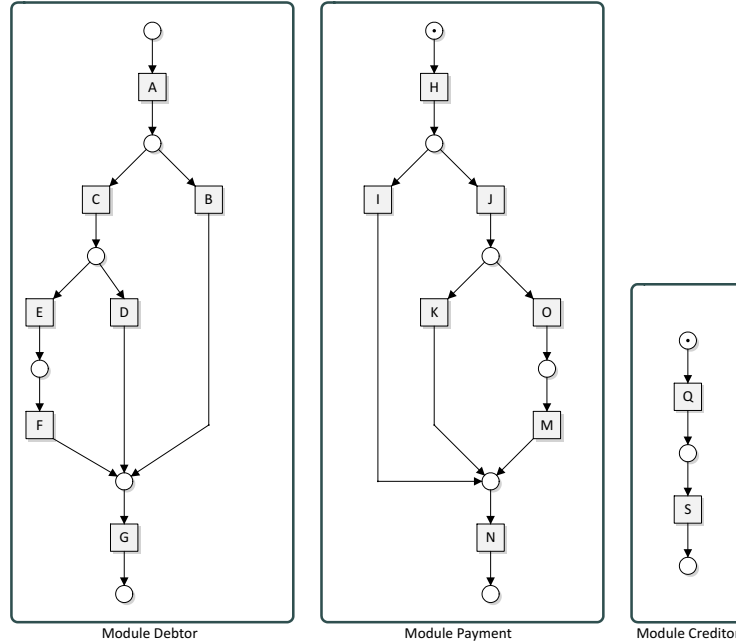


Fig. 6. Module nets generated for the running example

similarly each output place of the refined transition with each of the transitions in the preset of the final place of the refining workflow. It is straight-forward to prove that if (1) the initial net is sound, (2) each input place of the refined transition is 1-bounded, i.e., it can contain at most one token, and (3) workflow net W is sound, then the refinement yields a sound result.

The result of refining each feature by its feature net is shown in Fig. 7. As features G and N have no feature net defining communication, these transitions are not refined.

To verify whether the resulting open Petri nets are a true representation of the system, one can compose the nets into a single Petri net, and execute each of the sequences of the event log of Tbl. 1 on the resulting model, which in this example is possible. Further analyzing the resulting model shows that it is even a sound model: the final marking, i.e., the marking with only a token in the final places of the modules, is a home marking.

6 conclusions

Within this paper, we discussed a method to automatically generate a functional architecture model from an event log together with a mapping of each feature to the module that offers that functionality. We showed how the information flows can be derived from the communication behavioral profile. This profile not only

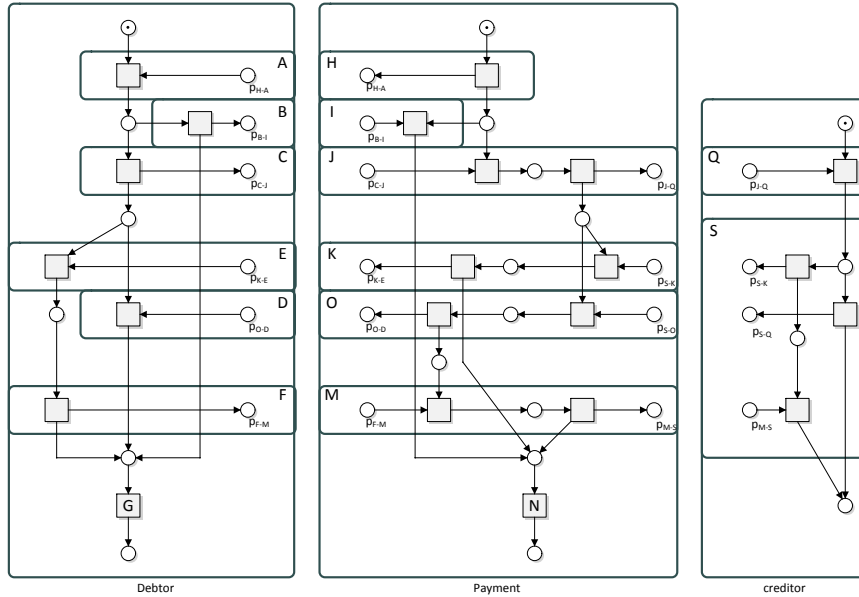


Fig. 7. Composed nets generated for the running example

identifies the information flow for the static structure of the functional architecture, but additionally offers sufficient information to construct the internal behavior for each of the features, and between the features within a module. Lastly, we showed how to compose feature and module nets into an open Petri net.

Discovering the interaction between different modules is not new. For example in service mining [1], where process mining techniques are used to derive the interaction between services. However, these techniques focus on creating a model of the interaction, whereas in this paper, we study the specification of the different modules. In [12], the authors discover the internal behavior of services based on the interaction between two services, guaranteeing deadlock freedom of the discovered service. In the setting of this paper, the exact interaction between modules is unknown, and needs to be discovered first.

Although the approach presented in this paper is promising, additional research is required. First, the current approach requires the event log to be complete, i.e., if the log grows, the successor relation should not change. Further, for the generation of the internal feature behavior, we assume that if the sending feature is present in the event log, it enables all possible events, which is possibly a too strict assumption that deserves further investigation.

The approach in this paper is very flexible, as we derive individual models for the feature and modules. For this, we apply standard process discovery algorithms returning sound workflow models. However, their composition in general

does not result in a sound system of asynchronously communicating systems. Further research is required to study the conditions under which this can be guaranteed. For this, we want to identify conditions which on the one hand result in correct models, and on the other hand have a positive effect on model quality as described by [8].

Not only does this approach provide useful insights for the software architect, we expect the approach applicable to business process management as well, as for the discovery of separate business processes, the Business Process Modelling and Notation offers the swimlane notion. We plan to implement the approach in the Process Mining toolkit ProM [15] to experiment and apply the approach on real-life examples.

References

- 1.
2. W. M. P. van der Aalst. Verification of Workflow Nets. In *Application and Theory of Petri Nets 1997*, volume 1248 of *Lecture Notes in Computer Science*, pages 407 – 426. Springer, Berlin, 1997.
3. W.M.P. van der Aalst. *Process Mining: Discovery, Conformance and Enhancement of Business Processes*. Springer, Berlin, 2011.
4. L. Bass, P. Clements, and R. Kazman. *Software Architecture in Practice*. Series in Software Engineering. Addison Wesley, Reading, MA, USA, 2012.
5. D. Bera, K. M. van Hee, and J. M. E. M. van der Werf. Designing weakly terminating ros systems. In *Applications and Theory of Petri Nets, (33th International Conference, Petri Nets 2012)*, volume 7347 of *Lecture Notes in Computer Science*, pages 328 – 347. Springer, Berlin, 2012.
6. S. Brinkkemper and S. Pachidi. Functional architecture modeling for the software product industry. In *ECISA 2010*, volume 6285 of *Lecture Notes in Computer Science*, pages 198 – 213. Springer, Berlin, 2010.
7. J.C.A.M. Buijs, B.F. van Dongen, and W.M.P. van der Aalst. A genetic algorithm for discovering process trees. In *Evolutionary Computation (CEC), 2012 IEEE Congress on*, pages 1–8, June 2012.
8. J.C.A.M. Buijs, B.F. van Dongen, and W.M.P. van der Aalst. On the Role of Fitness, Precision, Generalization and Simplicity in Process Discovery. In *On the Move to Meaningful Internet Systems: OTM 2012*, volume 7565 of *Lecture Notes in Computer Science*, pages 305–322. Springer, Berlin, 2012.
9. K.M. van Hee, N. Sidorova, and J.M.E.M. van der Werf. When can we trust a third party? In *Transactions on Petri Nets and Other Models of Concurrency VIII*, volume 8100 of *Lecture Notes in Computer Science*, pages 106–122. Springer, Berlin, 2013.
10. S.J.J. Leemans, D. Fahland, and W.M.P. van der Aalst. Discovering block-structured process models from event logs - a constructive approach. In *Application and Theory of Petri Nets and Concurrency*, volume 7927 of *Lecture Notes in Computer Science*, pages 311–329. Springer, Berlin, 2013.
11. G. Lucassen, J.M.E.M. van der Werf, and S. Brinkkemper. Alignment of software product management and software architecture with discussion models. In *IWSPM 2014*, pages 21–30. IEEE, Aug 2014.

12. R. Müeller, C. Stahl, W.M.P. van der Aalst, and M Westergaard. Service discovery from observed behavior while guaranteeing deadlock freedom in collaborations. In *Service-Oriented Computing*, volume 8274 of *Lecture Notes in Computer Science*, pages 358–373. Springer, Berlin, 2013.
13. W. Reisig. *Petri Nets: An Introduction*, volume 4 of *Monographs in Theoretical Computer Science: An EATCS Series*. Springer, Berlin, 1985.
14. R.N. Taylor, N. Medvidovic, and E.M. Dashofy. *Software Architecture: Foundations, Theory, and Practice*. John Wiley & Sons, 2010.
15. H.M.W. Verbeek, J.C.A.M. Buijs, B.F. van Dongen, and W.M.P. van der Aalst. XES, XESame, and ProM 6. In *Information System Evolution*, volume 72 of *Lecture Notes in Business Information Processing*, pages 60–75. Springer, Berlin, 2011.
16. M. Weidlich and J.M.E.M. van der Werf. On profiles and footprints – relational semantics for petri nets. In *Applications and Theory of Petri Nets, (33th International Conference, Petri Nets 2012)*, volume 7347 of *Lecture Notes in Computer Science*, pages 148 – 167. Springer, Berlin, 2012.