

INTERFACE MODELING FOR QUALITY AND RESOURCE MANAGEMENT *

MARTIJN HENDRIKS ^{a,b}, MARC GEILEN ^a, KEES GOOSSENS ^a, ROB DE JONG ^c,
AND TWAN BASTEN ^{a,b}

^a Eindhoven University of Technology, Eindhoven, The Netherlands

^b ESI (TNO), Eindhoven, The Netherlands

^c Philips Medical Systems International BV, Best, The Netherlands

ABSTRACT. We develop an interface-modeling framework for quality and resource management that captures configurable working points of hardware and software components in terms of functionality, resource usage and provision, and quality indicators such as performance and energy consumption. We base these aspects on partially-ordered sets to capture quality levels, budget sizes, and functional compatibility. This makes the framework widely applicable and domain independent (although we aim for embedded and cyber-physical systems). The framework paves the way for dynamic (re-)configuration and multi-objective optimization of component-based systems for quality- and resource-management purposes.

1. INTRODUCTION

Component frameworks enable the construction of systems from a set of loosely coupled, independent components through composition mechanisms. A component framework emphasizes abstraction, separation of concerns, and reuse. It has inherently loose coupling between components, and explicit and visible dependencies. This eases, for instance, dynamic reconfiguration and thereby adaptability. Component-based models of computation can be classified into component models and interface models [dAH01b]. A component model specifies how a component behaves in an arbitrary environment. An interface model specifies how a component can be used, e.g., for composition, and restricts the environment: it makes input assumptions and provides output guarantees for the case when the input assumptions are met by the environment. We present an interface model for (dynamic) quality and resource management (QRM) of component-based systems. Each component in our framework has one or more *configurations* that may be parameterized. A configuration

Key words and phrases: Constraint programming, multi-objective combinatorial optimization, Pareto algebra, embedded and cyber-physical systems, quality of service.

* The work described in this paper was supported from the ECSEL Joint Undertaking (JU) under grant agreement No 783162 (FitOpTiVis). The JU receives support from the European Union's Horizon 2020 research and innovation programme and Netherlands, Czech Republic, Finland, Spain, and Italy.

has an *input*, an *output*, a *required budget*, a *provided budget*, and a *quality*. Parameters capture the configurable working points of the component. External actors, the user or a quality and resource manager for instance, can use the parameters to partially control configuration of components for, e.g., dynamic reconfiguration of the system. The input and output model the functionality of a particular configuration of the component, to the extent relevant for QRM. For instance, a component configuration can have a video stream as input, and for each frame a set of image features as output. The types of extracted features may be adjustable via the parameters. The required and provided budgets model the services that a component configuration provides or requires. A component configuration may, for instance, require, or conversely, offer, a volatile storage capacity of 64 MB. Finally, the quality of a component configuration refers to the properties that we want to optimize, for instance, latency, throughput, energy consumption, cost. Qualities, budgets, and outputs typically depend on inputs and parameter settings. The five parts and the parameters form a *quality-and resource-management interface* that can be used for dynamic reconfiguration of the system to ensure functionality and a certain quality of service. We use concepts from the Pareto Algebra framework [GBTO07] to mathematically define our components and their composition.

Related work. Design based on components is a well-known approach to tackle the increasing complexity of today’s systems. The general idea is to use pre-defined components with well-specified behaviors and interfaces as interoperable building blocks. These components can be composed to create larger components with more complex functionality. Well-known component models in the software domain are the Component Object Model (COM) [Box98], the Common Object Request Broker Architecture (CORBA) [OMG12], JavaBeans [SM97], and Open Services Gateway Initiative (OSGi) [OA18]. These are all generally applicable and do not depend on the application domain. An example of a more specialized software component model is the KOALA framework [vOvdLKM00]. It is aimed at the software development of a product family of high-end television sets, with an emphasis on dealing with variability. Platforms that aim at predictability and composability are, for instance, the CompSOC platform [GKN⁺17] and PRET machines [EL07]. CompSOC aims to reduce system complexity through composability of applications and through (timing) predictability. PRET machines also focus on providing solid timing predictability for processors, thereby making the timing of the composition of programs predictable. Examples of component frameworks with tool support for cyber-physical systems are the Behavior Interaction Priorities (BIP) framework [BBS06] and the Ptolemy framework [EJL⁺03]. Other conceptual component models are, for instance, the I/O automata [LT87] and timed I/O automata [KLSV03] frameworks, which provide formal notions of components, composition, and abstraction.

Interface models for components have been introduced by [dAH01b]. An interface model of a component specifies what the component expects of the environment and what it can provide. This supports compositional refinement and therefore component-based design. Using this principle, interface automata capture temporal assumptions about how a component is used and how it uses other components [dAH01a]. The authors of, e.g., [WT05, HM06, KSHS17, GTW11] use interface-based design concepts for real-time aspects of systems. Contract-based design, e.g., [BCF⁺08, SVDP12], uses similar ideas. It specifies contracts of components as a combination of assumptions that the component makes about the environment and the guarantees that it then can give.

All work mentioned above has the scope of software development or system design. In contrast, we position our work in the scope of system operation with QRM as focus. In [APA⁺16], a middleware solution is presented that builds on a general-purpose component model. The middleware supports dynamic reconfiguration of components based on their quality-of-service demands. The quality-of-service parameters are fixed to period, deadline and priority, and each deployed component has a known worst-case execution time. The view presented in [APA⁺16] of components as sets of instances, each with its own quality-of-service parameters, is similar to our view. A runtime resource manager based on OpenCL is introduced in [LMFS13]. An extended compiler generates OpenCL kernels of the same code for various computational resources (CPU, GPU, FPGA, etc.) including resource-usage information. The runtime resource manager uses this to dynamically adapt the system if needed. The authors of [GBE10] informally introduce a component framework for quality-of-service management. They distinguish functional requirements, e.g., the required input format of a video-processing component, from resource requirements, e.g., the amount of memory needed. This is similar to our view in which we also make this distinction explicit. Finally, [SKB⁺15] presents a resource-management framework to create and manage multi-processor partitions at runtime, which provides a high level of (timing) predictability. The resource model is a simple instance of our model.

Contribution. Our framework uses partially-ordered sets as the interface model of components and defines composition in terms of operations from the Pareto Algebra [GBTO07]. We show that the framework supports a notion of refinement called *dominance*. This natively supports domain-independent quality and resource management by a formal semantics and compositional multi-objective optimization. Our work is on a formal, conceptual level and this distinguishes it from existing component-based approaches within the quality- and resource-management scope, e.g., [APA⁺16, LMFS13, GBE10], as these are concretely tailored to some domain or lack a formal semantics. The generality of our approach enables application to many domains, but requires further specialization before it can be applied.

Outline. We use a motivating example of a video-processing system from the healthcare domain throughout this paper. This example is introduced in Sec. 2. Sec. 3 summarizes the main results from [GBTO07] on Pareto Algebra that form the foundation of our interface model for QRM. In Sec. 4 we introduce our interface model, and in Sec. 5 we define composition of interfaces based on Pareto-algebraic operations. This is then detailed in Sec. 6 and applied to the example in Sec. 7. In Sec. 8, we introduce the QRML component language and toolset for the specification and analysis of quality- and resource-management problems and explain the semantics in terms of the QRM interface model developed in this paper. Finally, Sec. 9 concludes.

2. MOTIVATING EXAMPLE – DESCRIPTION

Example 2.1. We consider a video-processing system from the healthcare domain. The system scales video streams in an operating theatre, for display on one or more displays. The system receives video streams with a certain resolution and rate. Examples of common resolutions are qHD (960x540 pixels), HD (1280x720 pixels), HD+ (1600x900 pixels) and FHD (1920x1080 pixels). We assume that a pixel is encoded with four bytes. The platform of the system consists of fiber links, to transport streams from their sources to processing resources, and hardware scalars that perform the processing. Each fiber has a capacity of

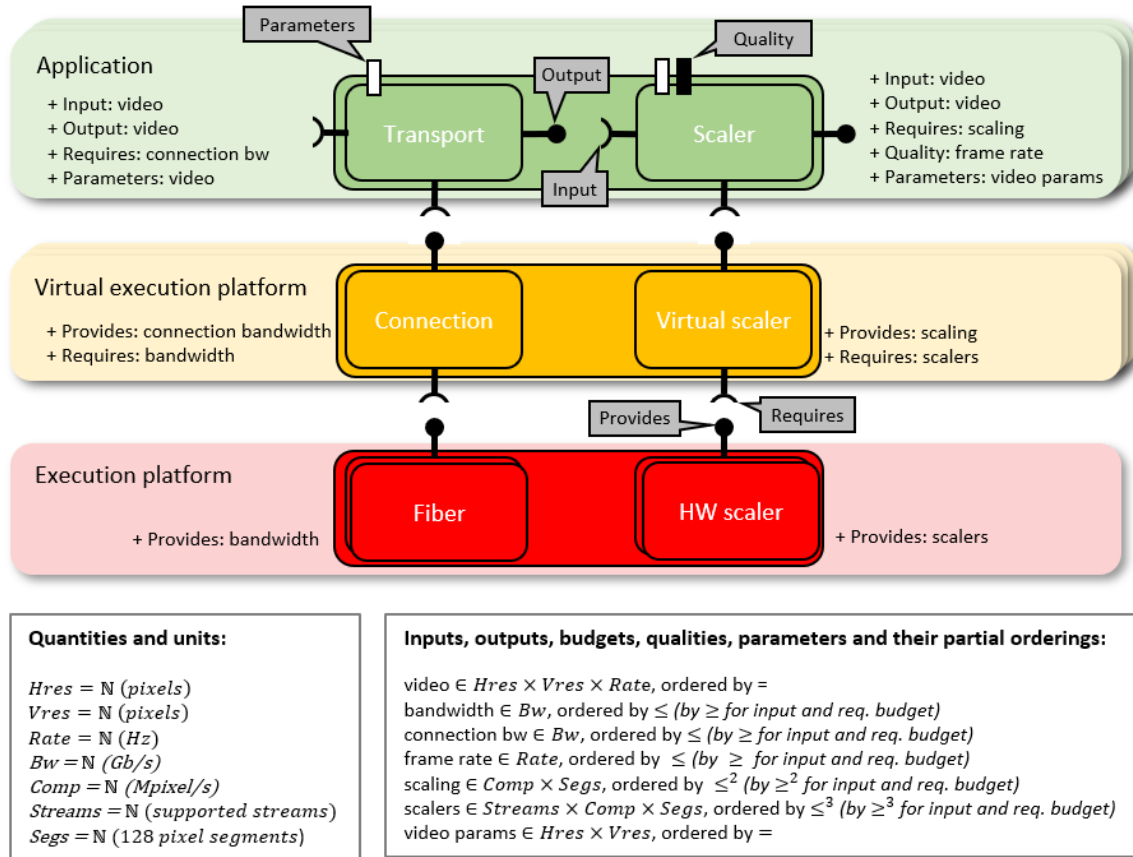


Figure 1: An example video-processing system with a single stream.

10 Gb/s. A single uncompressed FHD stream at 60 Hz takes approximately 4 Gb/s, so a fiber can transport at most two such streams. A scaler can downscale a stream to a given output resolution. Each scaler can process up to 300 Mpixels/s, and it can handle up to four streams. Depending on the resolution of a stream, a scaler requires a certain amount of memory to buffer a single horizontal line of that stream. Each scaler has access to only a limited amount of memory that thus must be shared by the streams that are mapped to that scaler. This memory is split into 32 segments of 128 pixels each. A scaler can reduce the frame rate of the outgoing stream to reduce its computational requirements. For instance, if a scaler does not have enough computational resources to process an incoming FHD stream at 60 Hz, then it can reduce the outgoing rate from 60 to 30 Hz by dropping half of the incoming video frames. This reduces the video quality at the benefit of being able to process more streams. The end user can also prioritize streams. The task of the quality and resource manager of this system is to map streams to scalers while respecting the resource constraints, and at the same time optimizing the frame rate of the outgoing streams given the end-user's priority. Dynamic reconfiguration takes place when streams are added or removed, or when the end user changes output resolutions or priorities.

Figure 1 shows an overview of the components in the system. Each stream in the system is modeled as an *application* that consists of a Transport component and a Scaler component. The Transport component has video both as input and output and requires

connection bandwidth that is determined by the rate and resolution of the video that is transported. This is a parameter of the Transport component that is set at runtime to match the transport configuration to the actual incoming stream. Video is modeled as a horizontal resolution in pixels, a vertical resolution in pixels, and a rate in Hz, each being a natural number. The connection bandwidth is the number of bits per second, again a natural number. The Scaler component also has video as input and output, and requires a scaling budget, which consists of computational capacity (Mpixels per second) and memory requirements (a number of 128 pixel segments for line buffering). The parameters of the Scaler specify working points for the resolution of the resulting video stream, and these can be set from the outside by, e.g., the end user or some application using the video system. The computational budget needed by the Scaler is determined by the input video resolution and the output video resolution and rate. The quality aspect of the Scaler is the frame rate of the outgoing video stream.

To have a feasible composition, the output of the Transport needs to *match* the input of the Scaler. We use partial-order relations to model this i/o matching, and also for the budget matching, i.e., matching a required budget with a provided budget. For instance, the $=$ relation could be used for *video*, as the Scaler can only guarantee its correct functioning if it exactly knows the input resolutions and rate. For connection bandwidth, for instance, we can use the \leq relation, as the transportation of a certain video can be guaranteed also when more bandwidth is given to the Transport component than strictly needed.

A Scaler component typically has multiple *configurations* that the end user or the quality and resource manager can partially select through the parameters of the component. A Transport component also has multiple configurations, where the specific configuration is determined by the environment since the type of the incoming stream is uncontrollable. For instance, suppose that we need to scale a 60 Hz FHD stream. (We abbreviate video-related types; for instance, we write FHD@60 for (1920, 1080, 60).) Then we have a Transport component with a configuration that requires a connection-bandwidth budget of 4 Gb/s, with a FHD@60 input and output video. The Scaler component has parameterized configurations. One configuration associated with output-resolution parameter HD has FHD@60 video input and HD@60 video output, requires a scaling budget of (171, 15), and has frame rate 60 as quality. Another configuration, also associated with output-resolution parameter HD, has FHD@60 video input and HD@30 video output, requires a scaling budget of (145, 15), and has frame rate 30 as quality. A third configuration with parameter HD has FHD@60 video input and HD@20 video output with scaling budget (136, 15), and quality 20. The latter two configurations trade quality for a smaller required budget.

The resource needs of an application, i.e., a Transport and Scaler pair in this example, are specified via a required private *virtual execution platform* (VEP) that abstracts from the underlying physical execution platform (for instance, real-time operating system and hardware). A VEP for our stream application consists of a Connection component that provides a connection bandwidth, and a Virtual Scaler component that provides scaling. The underlying hardware forms the *execution platform*, and consists of Fiber components and HW scaler components. A single Fiber provides bandwidth in Gb/s, which is modeled as a natural number. Each Scaler provides a budget consisting of three parts: (i) the number of streams that can still be processed by this Scaler (limited to four), (ii) the free amount of memory (at most 32 segments), and (iii) the remaining computational power (at most 300 Mpixels/s), captured by the tuple (4, 32, 300).

Again, suppose that a 60 Hz FHD stream needs to be scaled. The Transport component requires a connection bandwidth budget of 4 Gb/s, so a Connection component is used that provides this budget. To realize this, it requires a bandwidth budget of 4 Gb/s from some Fiber component. The remaining 6 Gb/s of that Fiber is available for other VEPs. In the same way, the required and provided budget of the Virtual scaler component can be calculated from the Scaler component (for a specific configuration of the Scaler component). A Virtual scaler component can be composed with any HW scaler component with matching budget.

In general, composition may lead to configurations that are not interesting because other configurations perform better in all aspects. Consider, for instance, two 60 Hz FHD streams that need to be scaled to HD. There are many possibilities to compose these with the hardware platform. A naive way is to compose both streams with a single HW scaler. This can be done if the output rate is set to 30 Hz. Another way is to compose each stream with its own Fiber and HW scaler. In that case, the output rates can be 60 Hz. The first option is *dominated* by the second option if the used and remaining platform resources are not important. Such a dominated combination can be ignored for further analysis and composition. That is, only the *Pareto-optimal* component combinations need to be considered.

The quality- and resource-management problem for this system is the following. Given a number of incoming video streams (applications) and given the required output resolution of each video stream (parameters), find mappings of these streams on a platform consisting of a number of Fibers and a number of HW scalers such that the output rates (qualities) are optimized according to some user-defined cost function (e.g., maximize the minimum output rate among all streams). In the following sections, we introduce the mathematics that we need to solve this problem. The application of the theory is demonstrated later in the paper, with a domain-specific solution for the quality- and resource-management problem of the video-processing system explained above.

3. PRELIMINARIES – PARETO ALGEBRA

This section briefly summarizes the results from Pareto Algebra [GBTO07] that we need for our interface-modeling framework. Prop. 3.25 is the only new part that we have added.

A partially-ordered set (poset) is a set Q with a partial-order relation \preceq_Q . In the remainder, we take shortcuts for readability and conciseness in the sense that we often write about a poset Q without mentioning the accompanying partial-order relation \preceq_Q explicitly. Furthermore, we sometimes omit the subscript Q from \preceq_Q when the ordering relation is clear from the context. The converse relation of \preceq is denoted by \preceq^{-1} and the strict (i.e., non-reflexive) version of the relation is denoted by \prec .

Definition 3.1 (Configuration Space). A configuration space \mathcal{S} is defined as the Cartesian product $Q_1 \times Q_2 \times \dots \times Q_n$ of a finite number of posets. A configuration is an element of a configuration space.

It is useful to have a notion whether a configuration is better than another configuration. This kind of information can help to choose between configurations in quality- and resource-management questions. The notion of *dominance* on a configuration space serves this purpose.

Definition 3.2 (Dominance). Let \mathcal{S} be a configuration space, and let $c = (q_1, \dots, q_n), c' = (q'_1, \dots, q'_n) \in \mathcal{S}$. We say that c is dominated by c' , or equivalently, that c' dominates c , denoted by $c \preceq c'$, if and only if $q_i \preceq_{Q_i} q'_i$ for all $1 \leq i \leq n$.

Note that we reverse the interpretation of \preceq with respect to [GBT07]. Observe that a configuration space with the dominance relation constructed from posets is itself again a poset. Hence, configuration spaces can serve as posets in other configuration spaces.

The notion of Pareto minimality expresses that a set of configurations only contains configurations that are relevant. This can be used to define the concepts of set dominance and Pareto equivalence of configuration sets.

Definition 3.3 (Pareto minimality). Let \mathcal{S} be a configuration space and let $C \subseteq \mathcal{S}$ be a set of configurations. We say that C is Pareto minimal if and only if $c \not\preceq c'$ for all $c, c' \in C$.

Definition 3.4 (Set dominance). Let \mathcal{S} be a configuration space and let $C, C' \subseteq \mathcal{S}$ be sets of configurations. We say that C' dominates C , denoted by $C \preceq C'$, if and only if for every $c \in C$ there is some $c' \in C'$ such that $c \preceq c'$.

Definition 3.5 (Pareto equivalence). Let \mathcal{S} be a configuration space and let $C, C' \subseteq \mathcal{S}$. We say that C is equivalent to C' , denoted by $C \equiv C'$ if and only if $C \preceq C' \wedge C' \preceq C$.

Pareto equivalence expresses the notion that neither of the sets have a configuration that cannot be at least matched by a configuration of the other set. The following theorem states that every well-ordered¹ set of configurations has a unique minimal equivalent, which is often called the Pareto frontier in other contexts. As a corollary of this theorem, every finite set of configurations has such a unique minimal equivalent.

Theorem 3.6. *If C is a set of configurations, and (C, \preceq^{-1}) is well ordered, then there is a unique Pareto-minimal set of configurations D such that $D \equiv C$.*

Corollary 3.7. *If C is a finite set of configurations, then there is a unique Pareto-minimal set of configurations D such that $D \equiv C$.*

Definition 3.8 (Minimization). Let \mathcal{S} be a configuration space, and let $C \in \mathcal{S}$ such that (C, \preceq^{-1}) is well ordered. We let $\min(C)$ denote the unique Pareto-minimal and Pareto-equivalent set of configurations.

Minimization of a set of configurations can be done with, for instance, the Simple Cull algorithm (see, e.g., [GB07]) with a worst-case complexity of $\mathcal{O}(n^2)$ where n is the number of configurations. Below we define some operations on configuration spaces.

Whether an operator on configuration spaces can safely use minimized operands is expressed through the notion of dominance preservation.

Definition 3.9 (Dominance preservation). Let \mathcal{S}, \mathcal{T} be configuration spaces, let $f : \mathcal{S} \rightarrow \mathcal{T}$ be an operator with n operands, and let $C_1, \dots, C_n, C'_1, \dots, C'_n \subseteq \mathcal{S}$. We say that f preserves dominance if and only if $(\forall_{1 \leq i \leq n} C_i \preceq C'_i) \Rightarrow f(C_1, \dots, C_n) \preceq f(C'_1, \dots, C'_n)$.

Corollary 3.10. *If f preserves dominance, then f can safely use Pareto-minimal operands, i.e., $f(C_1, \dots, C_n) \equiv f(\min(C_1), \dots, \min(C_n))$.*

Whether the result of an operator is already minimal is expressed by a notion of completeness that is formalized below.

¹A poset (Q, \sqsubseteq) is well ordered if every chain contains a smallest element.

Definition 3.11 (Minimality preservation). Let \mathcal{S}, \mathcal{T} be configuration spaces, let $f : \mathcal{S} \rightarrow \mathcal{T}$ be an operator with n operands, and let $C_1, \dots, C_n \subseteq \mathcal{S}$. We say that f *preserves minimality* if and only if $\min(f(C_1, \dots, C_n)) = f(\min(C_1), \dots, \min(C_n))$.

Pareto-minimal sets of configurations are sufficient and necessary for the optimization of arbitrary cost functions which behave monotonically.

Definition 3.12 (Derivation²). Let \mathcal{S} be a configuration space, let Q be a poset and let $f : \mathcal{S} \rightarrow Q$. We say that f is an \preceq -derivation (increasing derivation) if and only if $c \preceq c'$ implies that $f(c) \preceq_Q f(c')$ for all $c, c' \in \mathcal{S}$. Similarly, we say that f is a \succeq -derivation (decreasing derivation) if and only if $c \preceq c'$ implies that $f(c) \succeq_Q f(c')$ for all $c, c' \in \mathcal{S}$.

Definition 3.13 (Cost function). Let \mathcal{S} be a configuration space, let $C \subseteq \mathcal{S}$, let Q be a poset and let $f : \mathcal{S} \rightarrow Q$. Then we say that f is a cost function if and only if it is an \preceq -derivation. We say that $c \in C$ is an optimum in C if and only if $f(c') \preceq f(c)$ for all $c' \in C$, and c is called optimal if and only if $f(c) \not\preceq f(c')$ for all $c' \in C$.

The following two theorems, taken from [GBT07], express (i) that dominated configurations are not needed to derive optimal configurations, and (ii) that every Pareto-minimal configuration in a configuration set is the unique optimum under at least some cost function.

Theorem 3.14 (Sufficiency of Pareto-minimal sets). *Let \mathcal{S} be a configuration space, let $c, c' \in \mathcal{S}$, and let f be a cost function. If $c' \preceq c$, then $f(c') \preceq f(c)$.*

Theorem 3.15 (Necessity of Pareto-minimal sets). *Let \mathcal{S} be a configuration space, let $C \subseteq \mathcal{S}$ such that C is Pareto minimal, and let $c \in C$. Then there exists a cost function f such that c is the unique optimum configuration.*

Next, four operations on configuration spaces are defined that all preserve dominance and all but one preserve minimality. We use these operations later on to define composition of components in our framework.

Definition 3.16 (Free product). Let \mathcal{S}_1 and \mathcal{S}_2 be configuration spaces, let $C_1 \subseteq \mathcal{S}_1$ and let $C_2 \subseteq \mathcal{S}_2$. The free product of C_1 and C_2 is the Cartesian product $C_1 \times C_2$ in the configuration space $\mathcal{S}_1 \times \mathcal{S}_2$.

Proposition 3.17. *The free product preserves dominance.*

Proposition 3.18. *The free product preserves minimality.*

Definition 3.19 (Safe constraint). Let \mathcal{S} be a configuration space. A constraint is a set $D \subseteq \mathcal{S}$. A *safe* constraint is a set $D \subseteq \mathcal{S}$ such that for all $c_1, c_2 \in \mathcal{S}$ holds that if $c_2 \preceq c_1$ and $c_2 \in D$, then $c_1 \in D$ (implying that dominating configurations cannot be excluded from a safe constraint). Application of a constraint D to a set of configurations $C \subseteq \mathcal{S}$ is defined as $C \cap D$.

Proposition 3.20. *Application of a safe constraint preserves dominance.*

Proposition 3.21. *Application of a safe constraint preserves minimality.*

Definition 3.22 (Function application). Let \mathcal{S} be a configuration space, let $C \subseteq \mathcal{S}$, let Q be a poset and let $f : \mathcal{S} \rightarrow Q$. Application of f to \mathcal{S} yields a new configuration space $\mathcal{S} \times Q$ and is defined as $f(q_1, \dots, q_n) = (q_1, \dots, q_n, f(q_1, \dots, q_n))$. We define $f(C) = \{f(c) \mid c \in C\}$.

²This is a small and straightforward extension of the concept of *monotonicity* from [GBT07].

Note that this definition of function application overloads the function name. Function name f is also used to denote the application of f to some configuration space.

Proposition 3.23. *Application of an \preceq -derivation preserves dominance.*

Proposition 3.24. *Application of an \preceq -derivation preserves minimality.*

An application of a function can be written as an application of a constraint on the free product with the function range. This is formalized by the following proposition.

Proposition 3.25. *Let \mathcal{S} be a configuration space, let Q be a poset, and let $f : \mathcal{S} \rightarrow Q$. Then some constraint D exists such that $f(C) = D \cap (C \times Q)$ (where $f(C)$ is function application as in Def. 3.22) for all sets of configurations $C \subseteq \mathcal{S}$.*

Proof. We let $D = \{(s_1, \dots, s_n, f(s_1, \dots, s_n)) \mid (s_1, \dots, s_n) \in \mathcal{S}\}$.

(\subseteq) Let $c = (s_1, \dots, s_n, f(s_1, \dots, s_n)) \in f(C)$. By definition, $(s_1, \dots, s_n) \in C \subseteq \mathcal{S}$ and $f(s_1, \dots, s_n) \in Q$. So clearly $c \in C \times Q$. By definition, we also have that $c \in D$ and therefore $c \in D \cap (C \times Q)$.

(\supseteq) Let $c = (s_1, \dots, s_n, f(s_1, \dots, s_n)) \in D \cap (C \times Q)$. Thus, $(s_1, \dots, s_n, f(s_1, \dots, s_n)) \in (C \times Q)$ and therefore $(s_1, \dots, s_n) \in C$. By definition, $(s_1, \dots, s_n, f(s_1, \dots, s_n)) \in f(C)$. \square

From this proposition, it follows that the application of any \preceq -derivation can also be written as the application of a constraint, constructed as in the proposition, to the free product of the input set and the range of the derivation. This does not need to be a safe constraint. However, since application of a \preceq -derivation preserves dominance and minimality, this particular, equivalent, constraint application also preserves dominance and minimality. An example is the function $f : \mathbb{N} \rightarrow \mathbb{N}$ (with \leq as order on \mathbb{N}) defined as $f(n) = 2$. Clearly, this is a \preceq -derivation, and has the constraint $D = \{(n, 2) \mid n \in \mathbb{N}\}$. Although this is not a safe constraint, we can conclude that it nevertheless preserves dominance and minimality if used in $D \cap (C \times \mathbb{N})$ for all $C \subseteq \mathbb{N}$.

The dual of Prop. 3.25, which states that for every application of $D \cap (C \times Q)$ on a $C \subseteq \mathcal{S}$ we can find a function $f : \mathcal{S} \rightarrow Q$ that accomplishes the same, is not true (not even for safe constraints). A counter example is $D = \emptyset$, which is safe, but which cannot be mimicked by some function since that function application can never result in an empty set if $C \neq \emptyset$. Also note that $D \cap (C \times Q)$ can give a result C' such that $|C| < |C'|$ (because of the use of the free product). Again, this cannot be achieved with the derivation-based way.

Definition 3.26 (Abstraction). Let $\mathcal{S} = Q_1 \times Q_2 \times \dots \times Q_n$ be a configuration space, and let $C \subseteq \mathcal{S}$. The k -abstraction is a function $\mathcal{S} \rightarrow Q_1 \times \dots \times Q_{k-1} \times Q_{k+1} \times \dots \times Q_n$ defined as $(q_1, q_2, \dots, q_n) \downarrow k = (q_1, \dots, q_{k-1}, q_{k+1}, \dots, q_n)$. We define $C \downarrow k = \{c \downarrow k \mid c \in C\}$.

Proposition 3.27. *Application of k -abstraction preserves dominance.*

Abstraction does not preserve minimality because an abstraction can remove the single dimension that made some configuration incomparable to any other configuration. For instance, take $\mathcal{S} = \mathbb{N} \times \mathbb{N}$ with \leq as ordering on \mathbb{N} . Then $C = \{(1, 2), (2, 1)\}$ is Pareto minimal. We have $C \downarrow 1 = \{2, 1\}$, which clearly is not Pareto minimal because $1 < 2$.

The following is not explicit in [GBTO07] but follows trivially. Permutation allows us to reorder the dimensions of a configuration space.

Definition 3.28 (Permutation). Let $\mathcal{S} = Q_1 \times Q_2 \times \dots \times Q_n$ be a configuration space, let $C \subseteq \mathcal{S}$ be a set of configurations, and let $\pi : \{1, 2, \dots, n\} \rightarrow \{1, 2, \dots, n\}$ be a bijection,

i.e., a permutation. We define $\pi(\mathcal{S}) = Q_{\pi(1)} \times Q_{\pi(2)} \times \dots \times Q_{\pi(n)}$ and also let π denote the function $\mathcal{S} \rightarrow \pi(\mathcal{S})$ defined as $\pi(q_1, q_2, \dots, q_n) = (q_{\pi(1)}, q_{\pi(2)}, \dots, q_{\pi(n)})$. Furthermore, we define $\pi(C) = \{\pi(c) \mid c \in C\}$.

Proposition 3.29. *Permutation preserves dominance.*

Proposition 3.30. *Permutation preserves minimality.*

Finally, we need one more concept from Pareto Algebra that we use in our quality- and resource-management specialization later on.

Definition 3.31 (Alternatives). Let \mathcal{S} be a configuration space, and let $C_1, C_2 \subseteq \mathcal{S}$. Then $C_1 \cup C_2$ is called the set of alternatives of C_1 and C_2 .

Proposition 3.32. *Alternatives preserve dominance.*

The alternatives operation does not preserve minimality. For instance, take $\mathcal{S} = \mathbb{N}$ with \leq as ordering on \mathbb{N} . Then both $C_1 = \{1\}$ and $C_2 = \{2\}$ are Pareto minimal. However, $C_1 \cup C_2 = \{1, 2\}$ clearly is not Pareto minimal because $1 < 2$.

4. AN INTERFACE FOR QUALITY AND RESOURCE MANAGEMENT

The information needed for quality- and resource-management purposes is captured in our framework by an interface model for components. A component has an input, an output, a required budget, a provided budget, a quality, and it may be parameterized. The input, required budget, and parameters model the input assumptions, and the output, provided budget and quality model the output guarantees of the component [dAH01b]. All six parts are modeled by posets. Each of these posets can be multi-dimensional, to capture cases with multiple inputs, outputs, resource requirements, etc. This is the *quality- and resource-management interface* of the component.

Definition 4.1 (QRM interface). The QRM interface of a component is a set of configurations from a six-dimensional configuration space $Q_i \times Q_o \times Q_r \times Q_p \times Q_q \times Q_x$. By convention, the first dimension Q_i models the input of the component, the second dimension Q_o models the output, the third dimension Q_r models the required budget, the fourth dimension Q_p models the provided budget, the fifth dimension Q_q models the quality, and the sixth dimension Q_x models the parameters. Different components can use different configuration spaces.

A component can thus have a QRM interface that has multiple configurations. Each configuration models a specific working point of the component with specific values for input, output, required budget, provided budget, quality and parameters. These values usually are the result of some kind of analysis that is applied to the technology that is used to implement the component. External actors, e.g., end users, can typically control selection of subsets of configurations through the parameters. The freedom that remains to choose the final configuration is left to the quality and resource manager for optimization.

Figure 2 shows the graphical format that we use for a single configuration c of the QRM interface $C \subseteq Q_i \times Q_o \times Q_r \times Q_p \times Q_q \times Q_x$ of some component. Depicting the complete interface with all configurations is often problematic as the values of input i , output o , required budget r , provided budget p , quality q , and parameters x differ between configurations. The interpretation of the partial-order relation on input is that if $i \preceq i'$,

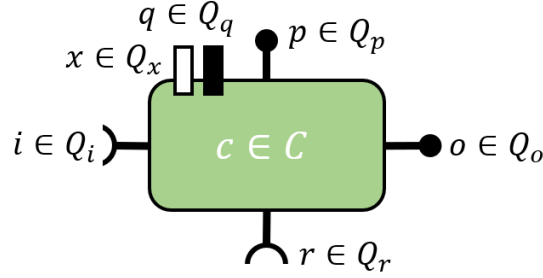


Figure 2: The interface for quality and resource management.

then i' is considered *weaker*³. For output, we interpret $o \preceq o'$ as that o' is *stronger* than o . For the required budget, we say that $r \preceq r'$ means that r' is *smaller* than r . For the provided budget, we say that $p \preceq p'$ means that p' is *larger* than p . The quality dimension is interpreted as follows: $q \preceq q'$ means that q' is *better* than q . Parameters are typically unordered; there are no better or worse parameter values. Whether a configuration is better or worse than, or incomparable to, another configuration is determined by the other five parts of the configurations.

In our framework, input- and required-budget specifications capture requirements, and output, provided budget and quality capture promises. Requiring more is worse than requiring less. We have diametrically opposed interpretations of \preceq for input and output, and also for required budget and provided budget, to be able to use the dominance relation of Def. 3.2. The dominating configuration thus has a weaker input requirement, a stronger output guarantee, a smaller required budget, a larger provided budget, a better quality and better, i.e., equal, parameters. This *alternating* approach to dominance is similar to the notion of alternating simulation that is used to define refinement for interface automata [dAH01a], and is also used, for instance, in contract-based design to define refinement between contracts [BCF⁺08]. Note that if we have a poset Q_{io} with partial order \preceq_i that models the input and we need a poset for modeling the output that has the same semantic domain, we can just take the dual poset, i.e., the same set Q_{io} with partial order \preceq_o where $x \preceq_o y$ if and only if $y \preceq_i x$ (i.e., $\preceq_o = \preceq_i^{-1}$). Clearly, the same goes for the required and provided budgets.

Example 4.2. Consider Ex. 2.1. The posets used to model the component parts are shown at the bottom of the figure. Now let us look at three configurations of the Scaler component, that scale FHS@60 input video to different resolutions and rates. Formally, these are the following⁴:

	input ($\mathbb{N}^3, =$)	output ($\mathbb{N}^3, =$)	required (\mathbb{N}^2, \geq^2)	provided ($\{\perp\}, =$)	quality (\mathbb{N}, \leq)	parameters ($\mathbb{N}^2, =$)
$s_1 =$	FHD@60,	HD+@60,	(201, 15),	\perp ,	60,	HD+
$s_2 =$	FHD@60,	HD+@30,	(160, 15),	\perp ,	30,	HD+
$s_3 =$	FHD@60,	HD@60,	(171, 15),	\perp ,	60,	HD

³For convenience we do not distinguish between the strict and non-strict versions for the description of the order relations.

⁴The Scaler component (and other components also) have many more configurations than mentioned in the examples. We only show some of these configurations to make our examples concise.

Here we use the *void poset* $\{\perp\}$ with partial order $=$ to model the provided budget of the Scaler. Note that the parameters poset uses equality as its ordering relation as parameters are unordered. Therefore, s_3 neither dominates nor is dominated by either s_1 or s_2 because it has a different parameter value. Because equality is also used for the output, s_1 neither dominates nor is dominated by s_2 . Now suppose that the Scaler has two other configurations that employ a different scaling algorithm that needs more buffering space:

$$\begin{aligned} s_4 &= \left(\text{FHD@60}, \text{HD+@60}, (201, 30), \perp, 60, \text{HD+} \right) \\ s_5 &= \left(\text{FHD@60}, \text{HD+@30}, (160, 30), \perp, 30, \text{HD+} \right) \end{aligned}$$

Then we have that $s_4 \preceq s_1$ and $s_5 \preceq s_2$ because the different scaling algorithm requires a larger budget but does not provide anything in terms of quality or output. Thus, $C = \{s_1, \dots, s_5\}$, is not Pareto minimal. Its Pareto-minimal and equivalent component is $D = \{s_1, s_2, s_3\}$, because $C \equiv D$, and D is Pareto minimal because $\{s_1, s_2, s_3\}$ forms an anti-chain with respect to the dominance relation. Adding the algorithm as an actual parameter would result in $\{s_1, \dots, s_5\}$ being Pareto minimal. Also the situation where s_4 and s_5 have an additional quality such as using less energy, would have this result.

The parameters can be used for several purposes. In this example, the parameters can be used to control the output resolution of the Scalers. We mentioned that parameters also can be used to differentiate between configurations that have different implementation algorithms. Such implementation alternatives may or may not be equal in terms of the input, output, budgets and quality. A related example is that they can be used to differentiate between different versions of components, where the version can be captured as an unordered quality to allow on-line version management.

A quality- and resource-management challenge arises when there are multiple components that implement the QRM interface. All these components must be combined in a proper way while optimizing some objective(s). In the next section, we therefore explain how Pareto-algebraic concepts can be used to create compositions of components that implement the QRM interface. These compositions can again be described by a QRM interface. The whole composition process is thus expressed in Pareto-algebraic terms, and the notion of Pareto minimality applies. This naturally gives the necessary tools for multi-objective optimization.

5. COMPOSITION OF QRM INTERFACES

We see composition as an *aggregation of alternatives*, and realize it with the Pareto-algebraic operations introduced in Sec. 3. The basic idea is that component interfaces can be realized with alternative options that are then aggregated in compositions as already illustrated in Fig. 1. The *alternatives pattern* and *aggregation pattern* are defined in the following subsections.

5.1. Alternatives pattern. We first generalize the binary alternatives operation of Def. 3.31 to an n -ary operation \cup^n . Note that the alternatives operation requires that the interfaces are from the same configuration space. Therefore, to apply the alternatives operation to interfaces $C_1 \subseteq \mathcal{S}_1, \dots, C_n \subseteq \mathcal{S}_n$, we must first *normalize* C_1, \dots, C_n to C'_1, \dots, C'_n respectively, such that $C'_1 \subseteq \mathcal{S}', \dots, C'_n \subseteq \mathcal{S}'$ for some configuration space \mathcal{S}' . This normalization is done

through proper derivations for each part of the interface, followed by abstraction and minimization. The pattern for taking the alternatives of n interfaces is the following:

$$\begin{aligned}
(1) \quad & \min \circ (\downarrow 1)^6 \circ f_x^1 \circ f_q^1 \circ f_p^1 \circ f_r^1 \circ f_o^1 \circ f_i^1 \\
& \vdots \\
(n) \quad & \min \circ (\downarrow 1)^6 \circ f_x^n \circ f_q^n \circ f_p^n \circ f_r^n \circ f_o^n \circ f_i^n \\
(n+1) \quad & \min \circ \cup^n
\end{aligned} \tag{5.1}$$

Here the f_x^1, \dots, f_i^n are the \leq -derivations that normalize the input, output, etcetera. Since these derivations add dimensions to the configuration space, we need to abstract from the original dimensions to get back to a six-dimensional configuration space with the QRM interpretation. The first step normalizes the first interface, the second step normalizes the second interface, etcetera. The $n+1$ -th step applies the alternatives operation to all normalized interfaces and minimizes the result.

The following proposition states that we can write the normalization as a constraint applied to a free product with the normalized configuration space.

Proposition 5.1. *Let \mathcal{S} be a configuration space, and let $f = f_x \circ f_q \circ f_p \circ f_r \circ f_o \circ f_i$ with $f_i : \mathcal{S} \rightarrow Q_i$, $f_o : \mathcal{S} \times Q_i \rightarrow Q_o$, etcetera. Then, some constraint D exists such that $f(C) = D \cap (C \times Q_i \times Q_o \times Q_r \times Q_p \times Q_q \times Q_x)$ for all $C \subseteq \mathcal{S}$.*

Proof. Let D_i , D_o , etcetera, be the constraints as defined in Prop. 3.25, for f_i , f_o , etcetera.

$$\begin{aligned}
& f_x \circ f_q \circ f_p \circ f_r \circ f_o \circ f_i(C) \\
& \quad = \text{Prop. 3.25} \\
& f_x \circ f_q \circ f_p \circ f_r \circ f_o(D_i \cap (C \times Q_i)) \\
& \quad = \text{Prop. 3.25} \\
& f_x \circ f_q \circ f_p \circ f_r \left(D_o \cap \left(\underbrace{D_i}_{X} \cap \underbrace{(C \times Q_i)}_Y \right) \times \underbrace{Q_o}_Z \right) \\
& \quad = (X \cap Y) \times Z = (X \times Z) \cap (Y \times Z) \\
& f_x \circ f_q \circ f_p \circ f_r \left(D_o \cap \left((D_i \times Q_o) \cap (C \times Q_i \times Q_o) \right) \right) \\
& \quad = X \cap (Y \cap Z) = (X \cap Y) \cap Z \\
& f_x \circ f_q \circ f_p \circ f_r \left((D_o \cap (D_i \times Q_o)) \cap (C \times Q_i \times Q_o) \right) \\
& \quad = D_o \cap (D_i \times Q_o) = D' \\
& f_x \circ f_q \circ f_p \circ f_r (D' \cap (C \times Q_i \times Q_o)) \\
& \quad = \dots \\
& D \cap (C \times Q_i \times Q_o \times Q_r \times Q_p \times Q_q \times Q_x) \quad \square
\end{aligned}$$

This proposition thus gives us a second way to accomplish normalization based on constraining the free product with the normalized configuration space $Q_i \times Q_o \times Q_r \times Q_p \times Q_q \times Q_x$. This results in the following constraint-based formulation of the alternatives pattern:

$$\begin{aligned}
(1) \quad & \min \circ (\downarrow 1)^6 \circ \cap D_1 \circ \times (\cdot, Q_i \times Q_o \times Q_r \times Q_p \times Q_q \times Q_x) \\
& \vdots \\
(n) \quad & \min \circ (\downarrow 1)^6 \circ \cap D_n \circ \times (\cdot, Q_i \times Q_o \times Q_r \times Q_p \times Q_q \times Q_x) \\
(n+1) \quad & \min \circ \cup^n
\end{aligned} \tag{5.2}$$

The proposition above shows that any normalization with derivations as in Eq. 5.1 can be written using constraints as in Eq. 5.2. The reverse is not true, i.e., the constraint-based

way is strictly more expressive than the derivation-based way. In line with the reasoning below Prop. 3.25, constraint-based normalization applied to a set of configurations C can result in the empty set (if the constraint is empty) and it may give a result C' such that $|C| < |C'|$. Both cannot be achieved with the derivation-based pattern.

Although the constraint-based alternatives pattern is more expressive, it can be easier to reason about the preservation of dominance and minimality - which is crucial in the Pareto Algebra framework - with the derivation-based pattern. The example below Prop. 3.25 shows that a simple \preceq -derivation cannot be mimicked with a safe constraint. As we see below however, this kind of derivation is useful for normalization. Only using safe constraints in Eq. 5.2 thus might limit the practical applicability. This consideration is the reason why we present both Eq. 5.1 and Eq. 5.2 as patterns for taking alternatives of components.

5.2. Aggregation pattern. Aggregation combines n constituent QRM interfaces into a new QRM interface. It consists of several steps and starts from a straightforward generalization of the free product in Def. 3.16 that combines n interfaces into a $6n$ -dimensional configuration space. Second, we apply constraints originating from, e.g., i/o and budget matching, and constraints on allowed parameter combinations to restrict the composition to feasible configurations. Third, we apply operations that express the semantics of the aggregation in the application domain. We assume that we have an aggregation operator for all six parts (input, output, etcetera) of the interface. Typically such an aggregation operator takes the parts of the constituent interfaces and creates a new part for the aggregation using a number of building-block functions, most notably domain-specific \preceq -derivations that can be applied using Def. 3.22. Fourth, after application of the proper derivations, we apply any relevant constraints on the resulting configurations. Fifth, we reconcile the result with our notion of QRM interface by abstraction of dimensions that represent the constituent interfaces. This results again in a six-dimensional configuration space of the aggregation with the usual interpretation of the dimensions. Finally, we apply Pareto minimization. All this is expressed in the following template:

$$\min \circ (\downarrow 1)^{6n} \circ \cap D_a \circ f_x \circ f_q \circ f_p \circ f_r \circ f_o \circ f_i \circ \cap D_c \circ \times^n \quad (5.3)$$

Here \times^n is the generalized free product, and f_i, \dots are the aggregation functions for input, output, etcetera. The constraints are split into two parts: D_c , applied directly after the free product on the constituent parts, and D_a , applied after the derivations that compute the interface values of the aggregation. Typically, i/o and budget matching, and parameter constraints are part of D_c . Domain-specific quality constraints, e.g., a constraint on a quality such as output rate, or constraints on the parameters of the composition can be part of D_a . Finally, $(\downarrow 1)^{6n}$ abstracts the dimensions of the constituent interfaces, and \min is the minimization operation.

Note that an actual implementation of this aggregation can be done much more efficiently than what is shown here. Typically, the application of constraints and the derivations and abstractions can be inlined during the creation of the free product. Since abstraction is always used in our composition and it does not preserve minimality, we always need to recalculate the set of Pareto-minimal configurations. This can be done using, e.g., the Simple Cull algorithm with quadratic worst-case complexity. Computation of the free product with the inline calculation of the constraints and derivations, can on its turn also be inlined in the Simple Cull algorithm, making it suitable for on-line operation [GB07].

Eq. 5.3 can be rewritten to the following form in which the derivations are replaced by a constraint in a similar way as is shown for the alternatives pattern above:

$$\min \circ (\downarrow 1)^{6n} \circ \cap D \circ \times (\cdot, Q_i \times Q_o \times Q_r \times Q_p \times Q_q \times Q_x) \circ \times^n \quad (5.4)$$

This pattern combines all the constraints in D_c and D_a with those originating from the aggregation functions in a single constraint D .

Composition of a number of components is accomplished by repeated application of the alternatives and aggregation patterns of Eq. 5.1, Eq. 5.2, Eq. 5.3, and Eq. 5.4: composition is an aggregation of alternatives. If we use dominance-preserving constraints in these operations, then the total composition function also preserves dominance since the other Pareto-algebraic operations used in the operation preserve dominance. Furthermore, since all operations end with a minimization, the total composition function also preserves minimality.

5.3. Refinement. The next theorem formalizes refinement in our setting. If component implementations are at least as good as is stated in their quality and resource management interfaces, then the composition of the implementations is at least as good as the composition based on the interfaces.

Theorem 5.2 (Refinement). *Let $C_1 \subseteq \mathcal{S}_1, \dots, C_n \subseteq \mathcal{S}_n$ and $C'_1 \subseteq \mathcal{S}_1, \dots, C'_n \subseteq \mathcal{S}_n$ be components, let $C_i \preceq C'_i$ for all $1 \leq i \leq n$, and let $f : 2^{\mathcal{S}_1} \times \dots \times 2^{\mathcal{S}_n} \rightarrow 2^{\mathcal{S}}$ be a function composition that only uses the basic Pareto-algebraic operations (free product, constraint application with safe constraints, derivation, abstraction, and alternatives, e.g., an aggregation of alternatives). Then, $f(C_1, \dots, C_n) \preceq f(C'_1, \dots, C'_n)$.*

Proof. By induction on the number of the basic Pareto-algebraic operations that we use in f . The base case for a single operation is proven by dominance preservation of the basic operators, see Props. 3.17, 3.20, 3.23, 3.27, 3.32. Now suppose that the theorem holds for $k - 1$ operations. Then we have composed an intermediate configuration set D from the set C_1, \dots, C_n , and we can compose an intermediate configuration set D' from the set C'_1, \dots, C'_n such that $D \preceq D'$. Applying the next operation to D and D' results in E and E' respectively. By dominance preservation of the composition operators and the induction hypothesis that $D \preceq D'$, we have that $E \preceq E'$. \square

A component declares a QRM interface consisting of a number of configurations C using Def. 4.1. This is an abstraction of the real configurations of the component, the implementation, that can also be expressed as a QRM interface C' , but which may not be exactly known. The assumption that we make, however, is that the implementation is at least as good as declared. That is to say that the implementation dominates the interface: $C \preceq C'$. Note that this allows the implementation to have more, but also less configurations than declared in the interface.

The theorem above is defined in terms of dominance of sets of configurations without taking a particular cost function into account. Now suppose that we have created a composition f based on the interface models and using a certain cost function we find an optimal configuration of this composition. I.e., we have $f(C_1, \dots, C_n) = C$ and the cost function selects $c \in C$ as optimum. We can trace for each interface which configuration was used to achieve this optimum: let $c_i \in C_i$ be the configurations used to obtain c . The question then is what happens when in an implementation, components and configurations are used that are refinements of the interface models, leading to the realization of the configuration c' instead of c ? According to the definition of set dominance, for every interface configuration

$c_i \in C_i$, there exists an implementation configuration c'_i in the implementation C'_i of the component, that is at least as good: $c_i \preceq c'_i$. From Th. 5.2, it follows that there exists a configuration c' in the composition of the implementations such that $c \preceq c'$. This thus means that an optimization done based on the interfaces is *safe* in the sense that the corresponding implementations yield a result that is at least as good for any arbitrary cost function.

The monotonic nature of our framework fits well with other formalisms with a monotonic nature that may be used to model and analyze the relation between the various quality and resource management aspects of the component configurations. An example is the synchronous dataflow (SDF) formalism [LM87] that is suitable for analyzing streaming applications such as the video-processing system of our motivating example. It is well-known that this formalism is monotone in the sense that reducing the execution times of actors⁵ will not result in a decrease of throughput. This is used in [GTW11] for compositional refinement that preserves key real-time performance metrics such as throughput and latency. We can also use the monotonicity in our setting. Suppose that some component can be modeled using SDF, and that its quality is throughput and its required budget relates directly to the execution times of the actors. SDF analysis gives concrete values for several component configurations, i.e., the analysis calculates the throughput for a number of required budget values (actor execution times). Now let us consider the composition with another component that provides the budget for the actor execution. The monotonicity of SDF in combination with Thm. 5.2 provides several robustness benefits. First, if the provided budget is larger than the required budget, then the throughput will not decrease because of SDF monotonicity. This means that the partial order on the actor execution budget can be a more flexible partial order than just equality. This eases implementation of the providing component: not exactly the required budget should be delivered, but a slightly larger budget (just to be safe) is also good according to Thm. 5.2. Second, usually the worst-case execution time of actors is used during SDF analysis. Because of monotonicity of SDF, an implementation will not have worse throughput. In terms of our framework, the component implementation dominates the interface that was used for composition. By Thm. 5.2, this is no problem and the result is at least as good as was predicted based on the interface.

6. DERIVATIONS AND CONSTRAINTS FOR COMPOSITION

In this section, we detail the use of Eq. 5.1 and Eq. 5.3 for composition, which we see as an aggregation of alternatives, using the running example. The use of the constraint-based way for composition, as defined by Eq. 5.2 and Eq. 5.4, is used and further discussed in Sec. 8, where we introduce a domain-specific language (DSL) for specifying the QRM view of a system with its semantics in terms of the framework developed in this paper.

6.1. Alternatives. The description of our running example in Sec. 2 does not contain alternatives. A variation of the case, however, also considers a software implementation of the scaler functionality that performs the processing. Remember that the hardware scaler provides a three-fold budget to model the supported number of streams, the computational power, and the available memory segments respectively. The software scaler can support a virtually unlimited number of streams, has virtually unlimited memory resources, but has

⁵An actor is part of the SDF formalism and models some (computational) task with a fixed execution time.

less computational power. It therefore only provides computation in its provided budget. In the following example, we show how we can create a component that models the choice between a hardware scaler and a software scaler.

Example 6.1 (Alternatives). The hardware scaler has configuration space $V \times V \times V \times \mathbb{N}^3 \times V \times V$ where $V = \{\perp\}$ with order = is the void poset. Its only configuration is the following:

$$h = (\perp, \perp, \perp, (4, 300, 32), \perp, \perp)$$

The software scaler has configuration space $V \times V \times V \times \mathbb{N} \times V \times V$ and also has a single configuration:

$$s = (\perp, \perp, \perp, 100, \perp, \perp)$$

To model the choice between these two components, we need to normalize them to a single configuration space. We choose the space $V \times V \times V \times (\mathbb{N} \cup \{\top\})^3 \times V \times V$. The special \top element is larger than any element of \mathbb{N} and thus models an unlimited provided budget.

Normalization of the provided budget of the hardware scaler is now modeled by the function $f_p^1 : \mathbb{N}^3 \rightarrow (\mathbb{N} \cup \{\top\})^3$ defined as $f_p^1(x, y, z) = (x, y, z)$ (i.e., it need not be changed). Normalization of the provided budget of the software scaler is modeled by the function $f_p^2 : \mathbb{N} \rightarrow (\mathbb{N} \cup \{\top\})^3$ defined as $f_p^2(y) = (\top, y, \top)$ (the unlimited budgets are added). The normalization functions for all parts but the provided budget are just copy functions $copy_i$ that copy dimension i (see below for a precise definition). These copy functions trivially are \preceq -derivations. Together, this gives the following instance of Eq. 5.1:

- (1) $min \circ (\downarrow 1)^6 \circ copy_6 \circ copy_5 \circ f_p^1 \circ copy_3 \circ copy_2 \circ copy_1$
- (2) $min \circ (\downarrow 1)^6 \circ copy_6 \circ copy_5 \circ f_p^2 \circ copy_3 \circ copy_2 \circ copy_1$
- (3) $min \circ \cup$

After application of the alternatives operation to $\{h\}$ and $\{s\}$, we obtain a general HWorSWscaler interface for the execution platform with two configurations:

$$\{(\perp, \perp, \perp, (4, 300, 32), \perp, \perp), (\perp, \perp, \perp, (\top, 100, \top), \perp, \perp)\}$$

6.2. Aggregation. As a first example of an aggregation, we show a free aggregation without any i/o or budget matching. This can, for example, be used to compose a platform from Fiber and HW scaler components. Some simple and generally useful \preceq -derivations that can be used to create the functions f_i , f_o , etcetera are defined next. First, a grouping derivation can be defined as follows that puts dimensions i to j in a single new multi-dimensional dimension. This new dimension uses the element-wise \preceq relation of its subdimensions (i.e., the dominance relation).

Definition 6.2 (Grouping). Let $\mathcal{S} = Q_1 \times \dots \times Q_n$ be a configuration space and let $1 \leq i \leq j \leq n$. We define $group_{i,j} : \mathcal{S} \rightarrow Q$ with $Q = Q_i \times \dots \times Q_j$ and $(p_i, \dots, p_j) \preceq_Q (q_i, \dots, q_j)$ if and only if $p_k \preceq_{Q_k} q_k$ for all $i \leq k \leq j$ as follows: $group_{i,j}(q_1, \dots, q_n) = (q_i, \dots, q_j)$.

Note that $group_{i,i}$ just copies a dimension. We abbreviate $group_{i,i}$ by $copy_i$.

Proposition 6.3. *Grouping is an \preceq -derivation.*

Proof. Directly from Def. 3.12. □

Next, ungrouping can help to access parts of multi-dimensional posets.

Definition 6.4 (Ungrouping). Let $\mathcal{S} = Q_1 \times \dots \times Q_n$ be a configuration space, let $1 \leq i \leq n$, let $Q_i = Q_i^1 \times \dots \times Q_i^m$ and let $1 \leq j \leq m$. Furthermore, let $(p_i^1, \dots, p_i^m) \preceq_{Q_i} (q_i^1, \dots, q_i^m)$ if and only if $p_i^k \preceq q_i^k$ for all $1 \leq k \leq m$. We define $ungroup_{i,j} : \mathcal{S} \rightarrow Q_i^j$ as $ungroup_{i,j}(q_1, \dots, q_{i-1}, (q_i^1, \dots, q_i^m), q_{i+1}, \dots, q_n) = q_i^j$.

Proposition 6.5. *Ungrouping is an \preceq -derivation.*

Proof. Directly from Def. 3.12. □

Finally, the void derivation adds a void value that can be used, e.g., to model empty component parts.

Definition 6.6 (Voiding). Let \mathcal{S} be a configuration space. We define $void : \mathcal{S} \rightarrow \{\perp\}$ as $void(c) = \perp$, with $=$ as the order on $\{\perp\}$.

Proposition 6.7. *Voiding is an \preceq -derivation.*

Proof. Directly from Def. 3.12. □

The next example shows how we apply all this for a simple aggregation.

Example 6.8 (Free aggregation). In this example, we aggregate a Fiber with a HW scaler, both introduced in Sec. 2. Since there is no connection made between input and output, or between provided and required budget, we call this a *free aggregation*. We assume that we have a single Fiber that provides 10 Gb/s bandwidth and a single HW scaler that provides a scalers budget of four streams, 300 Mpixels/s processing, and 32 memory segments of 128 pixels each. Both components have a single configuration, and only a non-void provided budget:

$$\begin{aligned} f_1 &= (\perp, \perp, \perp, 10, \perp, \perp) \\ h_1 &= (\perp, \perp, \perp, (4, 300, 32), \perp, \perp) \end{aligned}$$

We then apply the free product to obtain:

$$\{f_1\} \times \{h_1\} = \{ \overbrace{(\perp, \perp, \perp, 10, \perp, \perp)}^{f_1} \overbrace{(\perp, \perp, \perp, (4, 300, 32), \perp, \perp)}^{h_1} \}$$

Next, we define the compositions for the six interface parts. The composition of input, output and required budget is modeled by the void derivation, because the Fiber and HW scaler do not have these parts. The composition of the provided budget is done by grouping the respective provided budgets. Finally, we again use the void derivation for the composition of quality and parameters. This gives the following configuration:

$$\overbrace{(\perp, \perp, \perp, 10, \perp, \perp)}^{f_1} \overbrace{(\perp, \perp, \perp, (4, 300, 32), \perp, \perp)}^{h_1}, \perp, \perp, \perp, \perp, \perp, (10, (4, 300, 32)), \perp, \perp$$

The final steps are to abstract from the individual component parts, i.e., the first twelve dimensions, and to minimize (which is trivial for a configuration set of size one). This yields the execution platform component with only a single configuration.

$$e_1 = (\perp, \perp, \perp, (10, (4, 300, 32)), \perp, \perp)$$

This aggregation can thus be written as the following instance of Eq. 5.3 (where constraints are not present, as mentioned before):

$$min \circ (\downarrow 1)^{12} \circ void \circ void \circ group_{4,10} \circ void \circ void \circ void \circ \times$$

Our next example shows a more complicated aggregation. We want to make sure that input and output match. We do this by application of a producer-consumer constraint, which is already introduced in [GBT007].

Definition 6.9 (Producer-consumer constraint). Let $\mathcal{S} = Q_1 \times Q_2 \times \dots \times Q_n$ be a configuration space, let $1 \leq p \neq c \leq n$ and let $f : Q_c \rightarrow Q_p$ be a \succeq -derivation. Then, $D_{p,c,f} = \{(q_1, q_2, \dots, q_n) \in \mathcal{S} \mid f(q_c) \preceq_{Q_p} q_p\}$ is a producer-consumer constraint.

Proposition 6.10. *A producer-consumer constraint is safe.*

Proof. Let $d = (q_1, \dots, q_n), d' = (q'_1, \dots, q'_n) \in \mathcal{S}$ and let $d \preceq d'$. Let us assume that $d \in D_{p,c,f}$. We have to show that $d' \in D_{p,c,f}$, which is to say that $f(q'_c) \preceq_{Q_p} q'_p$. We have that $f(q_c) \preceq_{Q_p} q_p$ and since $d \preceq d'$ also that $q_c \preceq q'_c$ and $q_p \preceq q'_p$. Since f is a \succeq -derivation, we have that $f(q_c) \succeq_{Q_p} f(q'_c)$. Thus, $f(q'_c) \preceq_{Q_p} f(q_c) \preceq_{Q_p} q_p \preceq_{Q_p} q'_p$. \square

Another generally applicable constraint lets us select an arbitrary subset of configurations based on an unordered dimension. This is typically used to select arbitrary parameter valuations.

Definition 6.11 (Subset constraint). Let $\mathcal{S} = Q_1 \times Q_2 \times \dots \times Q_n$ be a configuration space, let $1 \leq i \leq n$, let $=$ be the partial order for Q_i , and let $X \subseteq Q_i$. The constraint $D_{i,X} = \{(q_1, q_2, \dots, q_n) \in \mathcal{S} \mid q_i \in X\}$ is called a subset constraint.

Proposition 6.12. *A subset constraint is safe.*

Proof. Let $c_1 = (p_1, \dots, p_n), c_2 = (q_1, \dots, q_n) \in \mathcal{S}$, let $c_2 \preceq c_1$, and let $c_2 \in D_{i,X}$. We have that $q_i \preceq p_i$ for all $1 \leq i \leq n$. We assumed that $c_2 \in D_{i,X}$ and thus $q_i \in X$. Because of the assumption on Q_i , we thus have that $p_i = q_i$, and, therefore, that $p_i \in X$. Thus, $(p_1, \dots, p_n) \in D_{i,X}$. \square

Example 6.13 (Horizontal aggregation). In this example, we aggregate a Transport with a Scaler. The key difference with the previous example is that we have to match the output of the Transport with the input of the Scaler. We call this a horizontal aggregation. In this example, we consider a Transport with two configurations: one for 60Hz FHD input, and one for 60Hz HD input. The required budget refers to the needed connection bandwidth in Gb/s.

$$\begin{aligned} t_1 &= (\text{HD@60}, \quad \text{HD@60}, \quad 2, \quad \perp, \quad \perp \quad \text{HD@60} \) \\ t_2 &= (\text{FHD@60}, \quad \text{FHD@60}, \quad 4, \quad \perp, \quad \perp \quad \text{FHD@60} \) \end{aligned}$$

For the Scaler, we look at the following configurations:

$$\begin{aligned} s_1 &= (\text{FHD@60}, \quad \text{HD+@60}, \quad (201, 15), \quad \perp, \quad 60, \quad \text{HD+} \) \\ s_2 &= (\text{FHD@60}, \quad \text{HD@30}, \quad (145, 15), \quad \perp, \quad 30, \quad \text{HD} \) \\ s_3 &= (\text{HD+@60}, \quad \text{HD@60}, \quad (135, 13), \quad \perp, \quad 60, \quad \text{HD} \) \end{aligned}$$

Taking the free component product results in a configuration set with six configurations (the left-hand side of these configurations is the Transport component):

$$\begin{aligned} a_1 &= (\text{HD@60}, \text{HD@60}, 2, \perp, \perp, \text{HD@60}, \quad \text{FHD@60}, \text{HD+@60}, (201, 15), \perp, 60, \text{HD+}) \\ a_2 &= (\text{HD@60}, \text{HD@60}, 2, \perp, \perp, \text{HD@60}, \quad \text{FHD@60}, \text{HD@30}, (145, 15), \perp, 30, \text{HD}) \\ a_3 &= (\text{HD@60}, \text{HD@60}, 2, \perp, \perp, \text{HD@60}, \quad \text{HD+@60}, \text{HD@60}, (135, 13), \perp, 60, \text{HD}) \\ a_4 &= (\text{FHD@60}, \text{FHD@60}, 4, \perp, \perp, \text{FHD@60}, \quad \text{FHD@60}, \text{HD+@60}, (201, 15), \perp, 60, \text{HD+}) \\ a_5 &= (\text{FHD@60}, \text{FHD@60}, 4, \perp, \perp, \text{FHD@60}, \quad \text{FHD@60}, \text{HD@30}, (145, 15), \perp, 30, \text{HD}) \\ a_6 &= (\text{FHD@60}, \text{FHD@60}, 4, \perp, \perp, \text{FHD@60}, \quad \text{HD+@60}, \text{HD@60}, (135, 13), \perp, 60, \text{HD}) \end{aligned}$$

The stream that needs to be processed is of type FHD@60, and we use a subset constraint on the Transport parameter, $D_{6,\{\text{FHD@60}\}}$, to select only the relevant transport configurations, which are a_4 , a_5 and a_6 . We also need to match the output of the Transport to the input of the Scaler. Since the output type of the Transport is equal to the input type of the Scaler, we use the identity function \mathbf{id} (which clearly is a \succeq -derivation) to create a producer-consumer constraint $D_{2,7,\mathbf{id}}$. Then we see that $a_6 \notin D_{2,7,\mathbf{id}}$, since output FHD@60 does not match input HD+@60 (with = as the ordering). This removes a_6 . Then we also apply a parameter constraint, $D_{12,\{\text{HD}\}}$, to select only those configurations with an HD output. That removes a_4 since it has HD+ output.

The derivations for the parts are the following. We use $copy_1$ for the input, since the input of the aggregation is the input of the Transport. We use $copy_8$ for the output, since the output of the aggregation is the output of the Scaler. This models that the output of the transport is consumed completely by the Scaler. We use $group_{3,9}$ for the required budget (collecting the required budgets of the two constituent interfaces), $void$ for the provided budget, $copy_{11}$ for the quality (indicating that the aggregation inherits its quality from the Scaler) and $void$ for the parameters. This last voiding models abstraction of the Transport and Scaler parameters as they have served their purpose and are not relevant anymore. Alternatively, one could decide that the aggregation copies the parameters from the Scaler, to allow configuration of the output of streams as a whole. Finally, we abstract from the first twelve dimensions which contain the individual component parts and minimize. This gives the following configuration:

$$a_5 = (\text{FHD@60}, \text{HD@30}, (4, (145, 15)), \perp, 30, \perp)$$

This aggregation formally is the following instance of Eq. 5.3:

$$\begin{aligned} \min \circ (\downarrow 1)^{12} \circ \text{void} \circ \text{copy}_{11} \circ \text{void} \circ \text{group}_{3,9} \circ \text{copy}_8 \circ \text{copy}_1 \circ \\ \cap D_{12,\{\text{HD}\}} \circ \cap D_{6,\{\text{FHD@60}\}} \circ \cap D_{2,7,\mathbf{id}} \circ \times \end{aligned}$$

This example shows how we can use parameters to select subsets of configurations, and how to use a producer-consumer constraint for i/o matching (budget matching works in a similar way). The output in the example is completely consumed. This, however, needs not be the case in general. Modeling broadcast, for instance, leaves the broadcasting output untouched in an aggregation. In general, we need non-trivial \preceq -derivations for the aggregation of the parts, for instance, to calculate the remaining budget when a required budget is strictly smaller than a provided budget. The remaining budget is then provided to other components by the aggregation. Another example is the aggregation of two image-processing components that both have latency as quality. For the aggregation, we want to add the latencies of the individual components. Real numbers can often be used to express inputs, outputs, budgets and qualities. The kind of arithmetic operations that we need for the proper \preceq -derivations are supported by the following definitions.

Definition 6.14 (Addition-Multiplication-Min-Max). Let $\mathcal{S} = Q_1 \times \dots \times Q_n$ be a configuration space, let $1 \leq i, j \leq n$, let $Q_i = \mathbb{R}$ and $Q_j = \mathbb{R}$, and let $\preceq_{Q_i} = \leq$ and $\preceq_{Q_j} = \leq$. We define $add_{i,j} : \mathcal{S} \rightarrow Q$ with $Q = \mathbb{R}$ and $\preceq_Q = \leq$ as $add_{i,j}(q_1, \dots, q_n) = q_i + q_j$. We define $mult$, min and max in a similar way.

Proposition 6.15. *Addition/multiplication/min/max are \preceq -derivations.*

Proof. Let $(q_1, \dots, q_n), (q'_1, \dots, q'_n) \in \mathcal{S}$ and let $(q_1, \dots, q_n) \preceq (q'_1, \dots, q'_n)$. We have to show that $f(q_i, q_j) \leq f(q'_i, q'_j)$ with f being addition, multiplication, min or max. This follows straightforwardly from the assumption that $q_i \leq q'_i$ and $q_j \leq q'_j$. \square

For subtraction and division, the situation is subtly different as the ordering on one of the posets needs to be reversed. Subtraction can be used, for instance, to compute the remaining provided budget when a budget provider is aggregated with a budget consumer. This, however, nicely matches with the alternating interpretation of input and output, and provided budget and required budget as we will see in the example below. In the following definition, Q_p is typically the producer (output or provided budget), and Q_c is the consumer (input or required budget).

Definition 6.16 (Subtraction-Division). Let $\mathcal{S} = Q_1 \times \dots \times Q_n$ be a configuration space, let $1 \leq p \neq c \leq n$, let $Q_p = Q_c = \mathbb{R}$, and let $\preceq_{Q_p} = \leq$ and $\preceq_{Q_c} = \geq$. We define $sub_{p,c} : \mathcal{S} \rightarrow Q$ with $Q = \mathbb{R}$ and $\preceq_Q = \leq$ as $sub_{p,c}(q_1, \dots, q_n) = q_p - q_c$. We define div in a similar way, but only for $Q_p = Q_c = \mathbb{R}^{\geq 0}$.

Proposition 6.17. *Subtraction/division are \preceq -derivations.*

Proof. Let $(q_1, \dots, q_n), (q'_1, \dots, q'_n) \in \mathcal{S}$ and let $(q_1, \dots, q_n) \preceq (q'_1, \dots, q'_n)$. We have to show that $f(q_p, q_c) \leq f(q'_p, q'_c)$ with f being subtraction or division. We have assumed that $q_p \preceq q'_p$, which means that $q_p \leq q'_p$. We also have assumed that $q_c \preceq q'_c$, which means that $q_c \geq q'_c$. Clearly, $q_p - q_c \leq q'_p - q'_c$, and (assuming non-negative numbers) $q_p/q_c \leq q'_p/q'_c$. \square

Example 6.18 (Vertical aggregation). In this example, we aggregate a Fiber component that provides 10 Gb/s *bandwidth* and a Connection component that requires 4 Gb/s *bandwidth* and provides 4 Gb/s *connection bandwidth*. We connect the provided budget of the Fiber with the required budget of the Connection and therefore call this a vertical aggregation.

$$\{f_1\} = \{(\perp, \perp, \perp, 10, \perp, \perp)\}$$

$$\{c_1\} = \{(\perp, \perp, 4, 4, \perp, \perp)\}$$

$$\{f_1\} \times \{c_1\} = \{(\perp, \perp, \perp, 10, \perp, \perp, \perp, \perp, 4, 4, \perp, \perp)\}$$

Application of the budget constraint $D_{4,9,\text{id}}$ does not filter out this configuration because $4 \leq 10$. Then, we apply *void* three times to model the input and output derivation, and the derivation for the required budget. We use $sub_{4,9}$ (on the natural numbers) to model the consumption of the budget and compute the remaining provided *bandwidth* budget. Note that the requirements that $Q_4 = Q_9 = \mathbb{N}$, $\preceq_{Q_4} = \leq$ and $\preceq_{Q_9} = \geq$ are met because of the alternating approach to modeling provided and required budgets. We group this remaining provided *bandwidth* with the provided *connection bandwidth* and remove the superfluous dimension. Finally, we apply *void* twice to model the composition of the quality and parameter parts. As last steps, we abstract from the first twelve dimensions and minimize. This gives the following configuration, capturing that the aggregate component provides a combination of 4 Gb/s *connection bandwidth* and 6 Gb/s raw *bandwidth*:

$$(\perp, \perp, \perp, (4, 6), \perp, \perp)$$

This aggregation formally is the following instance of Eq. 5.3:

$$min \circ (\downarrow 1)^{12} \circ void^2 \circ (\downarrow 16 \circ group_{10,16} \circ sub_{4,9}) \circ void^3 \circ \cap D_{4,9,\text{id}} \circ \times$$

So far, we have given a number of operations that we can use for aggregation. The free product combines components and creates all possible combinations of the individual configurations. Abstraction can be used to remove dimensions that are not relevant anymore. Function application of \preceq -derivations defined above gives us operations that can be used to model the aggregation of input, output, etcetera. Finally, the producer-consumer constraint can be used to match input with output or required budget with provided budget, and the subset constraint can be used to select an arbitrary subset of configurations based on an unordered dimension (e.g., the parameter dimension). As shown in Ex. 6.8, Ex. 6.13 and Ex. 6.18, aggregation functions can be constructed from these basic building blocks using the template of Eq. 5.3.

6.3. Aggregation templates. The examples in the previous section have shown three templates that can be used for aggregation: free aggregation, horizontal aggregation and vertical aggregation. These templates can be realized by using the general aggregation pattern of Eq. 5.3 with proper derivations to implement f_i etc., and with constraints for i/o and budget matching, parameter selection, and domain-specific filtering using the quality dimension. Next, we show how we can realize these templates in terms of generalized addition and subtraction operators.

Let us assume that we know how to add and subtract, i.e., we have functions $+$ and $-$ with signature $Q \times Q \rightarrow Q$ with the following properties for every poset Q :

$$\begin{aligned} a \preceq_Q c \wedge b \preceq_Q d &\Rightarrow a +_Q b \preceq_Q c +_Q d \\ a \preceq_Q c \wedge b \succeq_Q d &\Rightarrow a -_Q b \preceq_Q c -_Q d \end{aligned}$$

These properties hold for the regular addition and subtraction of numbers. These operations give rise to \preceq -derivations. Poset addition, defined below, takes care of adding values from identical posets, and also of values from different posets by grouping them together in a tuple.

Definition 6.19 (Poset addition). Let $\mathcal{S} = Q_1 \times \dots \times Q_n$ be a configuration space, and let $1 \leq i, j \leq n$. We define $+_{i,j} : \mathcal{S} \rightarrow Q$ as follows: $+_{i,j}(q_1, \dots, q_n) = q_i +_{Q_i} q_j$ with $\preceq_Q = \preceq_{Q_i}$ if $Q_i = Q_j$ and $\preceq_{Q_i} = \preceq_{Q_j}$ and otherwise $+_{i,j}(q_1, \dots, q_n) = (q_i, q_j)$ with $(q_i, q_j) \preceq_Q (q'_i, q'_j)$ if and only if $q_i \preceq_{Q_i} q'_i$ and $q_j \preceq_{Q_j} q'_j$.

Proposition 6.20. *Poset addition is an \preceq -derivation.*

Proof. Let $(q_1, \dots, q_n), (q'_1, \dots, q'_n) \in \mathcal{S}$ and let $(q_1, \dots, q_n) \preceq (q'_1, \dots, q'_n)$. First, assume that $Q_i = Q_j$ and $\preceq_{Q_i} = \preceq_{Q_j}$. Then, $+_{i,j}(q_1, \dots, q_n) = q_i +_{Q_i} q_j$ and $+_{i,j}(q'_1, \dots, q'_n) = q'_i +_{Q_i} q'_j$. From the assumption that $q_i \preceq_{Q_i} q'_i$ and $q_j \preceq_{Q_i} q'_j$, it follows with the definition of $+_{Q_i}$ that $q_i +_{Q_i} q_j \preceq_{Q_i} q'_i +_{Q_i} q'_j$. Second, assume that $Q_i \neq Q_j$ or $\preceq_{Q_i} \neq \preceq_{Q_j}$. Then $+_{i,j}(q_1, \dots, q_n) = (q_i, q_j)$ and $+_{i,j}(q'_1, \dots, q'_n) = (q'_i, q'_j)$. Clearly, $(q_i, q_j) \preceq_Q (q'_i, q'_j)$ since $q_i \preceq_{Q_i} q'_i$ and $q_j \preceq_{Q_j} q'_j$. \square

The case for subtraction is more complicated due to the alternating interpretation but similar to the situations of the producer-consumer constraint of Def. 6.9 and the subtraction-division of Def. 6.16.

Definition 6.21 (Poset subtraction). Let $\mathcal{S} = Q_1 \times \dots \times Q_n$ be a configuration space, let $1 \leq i, j \leq n$, let $Q_i = Q_j$, and let $\preceq_Q = \preceq_{Q_i}^{-1}$. We define $-_{i,j} : \mathcal{S} \rightarrow Q$ as follows: $-_{i,j}(q_1, \dots, q_n) = q_i -_{Q_i} q_j$ with $\preceq_Q = \preceq_{Q_i}$.

Proposition 6.22. *Poset subtraction is an \preceq -derivation.*

Proof. Let $(q_1, \dots, q_n), (q'_1, \dots, q'_n) \in \mathcal{S}$ and let $(q_1, \dots, q_n) \preceq (q'_1, \dots, q'_n)$. Then we have that $-_{i,j}(q_1, \dots, q_n) = q_i -_{Q_i} q_j$ and $-_{i,j}(q'_1, \dots, q'_n) = q'_i -_{Q_i} q'_j$. We must show that $q_i -_{Q_i} q_j \preceq_{Q_i} q'_i -_{Q_i} q'_j$. From our assumptions, we have that $q_i \preceq_{Q_i} q'_i$ and $q_j \preceq_{Q_j} q'_j$. Since $Q_i = Q_j$ and $\preceq_{Q_i} = \preceq_{Q_j}^{-1}$, we have that $q_j \succeq_{Q_i} q'_j$. With the definition of $-_{Q_i}$, this gives that $q_i -_{Q_i} q_j \preceq_{Q_i} q'_i -_{Q_i} q'_j$. \square

With these operations on posets, we can then define free, horizontal and vertical aggregation as follows (where the finishing steps, minimization and abstraction of the twelve dimensions and filtering on parameters or quality constraints, are omitted):

$$\begin{aligned} & \parallel \stackrel{\text{def}}{=} && +_{6,12} \circ +_{5,11} \circ +_{4,10} \circ +_{3,9} \circ +_{2,8} \circ +_{1,7} \circ \times \\ \Rightarrow & \stackrel{\text{def}}{=} && +_{6,12} \circ +_{5,11} \circ +_{4,10} \circ +_{3,9} \circ (\downarrow 14 \circ +_{8,14} \circ -_{2,7}) \circ \text{copy}_1 \circ \cap D_{2,7,\mathbf{id}} \circ \times \\ \Uparrow & \stackrel{\text{def}}{=} && +_{6,12} \circ +_{5,11} \circ (\downarrow 16 \circ +_{10,16} \circ -_{4,9}) \circ \text{copy}_3 \circ +_{2,8} \circ +_{1,7} \circ \cap D_{4,9,\mathbf{id}} \circ \times \end{aligned}$$

Horizontal aggregation is only defined if the output type of the first component is equal to the input type of the second component, and if the orderings on input and output are reversed (due to the requirement of the poset subtraction). Our definition is such that the input of the second component is completely covered by the output of the first component (the producer-consumer constraint). What remains of the output of the first component is added to the output of the second component resulting in the output of the aggregation. The definition of vertical aggregation is similar, only the roles of input and output are played by required and provided budget, respectively. The three aggregation templates are shown schematically in Fig. 3, Fig. 4 and Fig. 5, respectively. Note that these are specific example templates of how we can define aggregation with the addition and subtraction operations on posets. For instance, the specific definitions for the consumption of output and provided budget is a choice. We can imagine that there are situations where variations of these aggregation templates are needed. These can then be defined with the flexible mathematical tools that we have presented.

When not everything is known at design time, which is the case in our example because the video streams that need processing are only known at runtime, an online quality and resource manager is needed. In the following section we develop such a manager which is fully tailored to the video-processing example and implemented in the Java programming language.

7. QUALITY AND RESOURCE MANAGEMENT FOR THE VIDEO-PROCESSING SYSTEM

In this section, we put all theory into practice and create a Java solution for the quality- and resource-management problem as described in Ex. 2.1. We emphasize that this is not a general approach to quality and resource management, but an example of how the mathematical tools that we have provided can be used in practice to solve a specific problem. The first step was to implement the posets that we need, e.g., *video*, *bandwidth*, etcetera (see Fig. 2.1). We capture the posets in the `Value` type (a Java interface), which provides ordering methods to compare the `Value` with another `Value`, and also the $+$ and $-$ operations defined in Sec. 6.3. We have an abstract `NatValue` implementation of this interface for natural numbers, and several extensions of `NatValue` to model the concrete types that we need, e.g.,

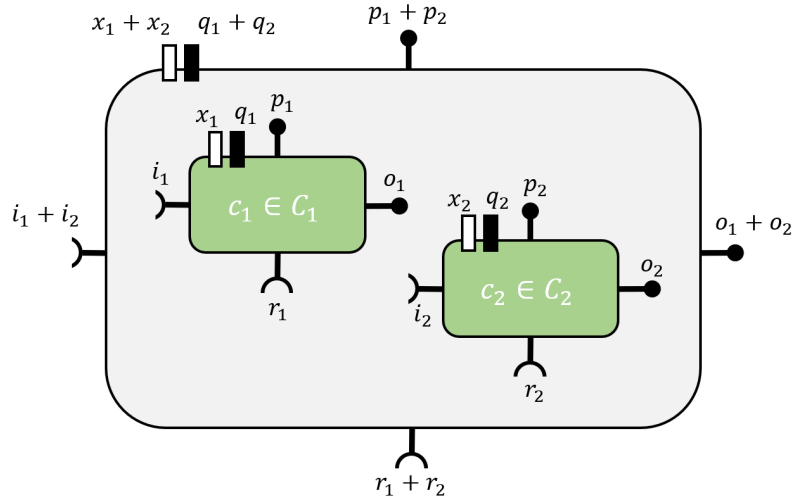


Figure 3: Abstract representation of free aggregation.

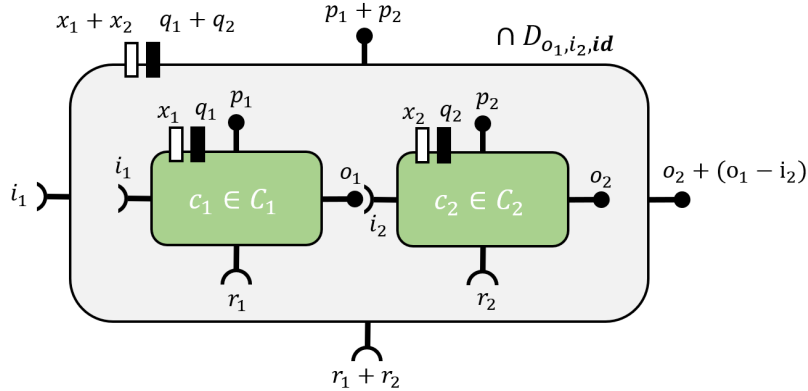


Figure 4: Abstract representation of horizontal aggregation.

bandwidth, *hres*, *vres*, etcetera. Furthermore, we have a `Tuple` implementation of the `Value` interface that realizes tuples of arbitrary length and nesting, with other `Value` instances at its leafs. We have several extensions of this `Tuple` type to model the multi-dimensional base types that we need, e.g., *video* and *scalers*.

As second step, we implemented the Pareto-algebraic machinery specialized to quality and resource management. We defined an interface `ConfigurationSet` that components can implement. It has a single method to get all configurations of the component, which are implementations of the `Configuration` interface. This type provides the methods to obtain `Value` instances for input, output, etcetera. Furthermore, we implemented all Pareto-algebraic operations that we need, e.g., minimization, abstraction, etcetera. We also implemented the derived functions that we need, e.g., the free, horizontal and vertical aggregation using the $+$ and $-$ operations on the posets, as described above.

As a third step, we created an application-specific QRM algorithm to map `Transport` and `Scaler` components for n streams on k `Fiber` and k `HW scaler` components using appropriate `VEP` components. We partitioned the platform resources into execution-platform components.

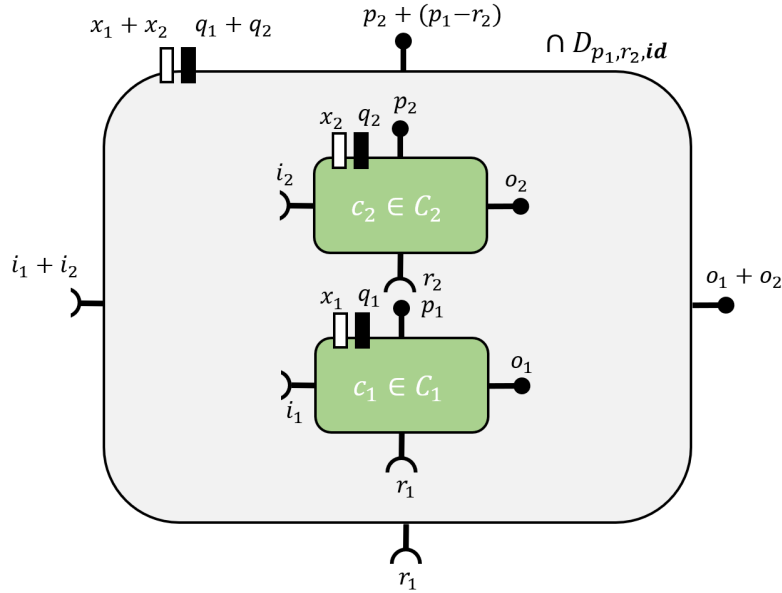


Figure 5: Abstract representation of vertical aggregation.

Each is a free aggregation of a Fiber and a HW scaler component. A stream is modeled by a single application component, which is the horizontal aggregation of a Transport component with a Scaler component, see Ex. 6.13. The configurations of the Scaler are determined as follows:

- ▷ The input can be FHD, HD+, HD, qHD (960x540 pixels) or nHD (640x360 pixels) with a rate of 90, 60, 30, 20, 15, 10, 6, or 5 Hz.
- ▷ The output can be as the input, provided that the resolution of the output is smaller than the resolution of the input, and the rate of the output divides the rate of the input.
- ▷ The required *scaling* budget consists of two parts. The *comp* part is calculated as the input pixel rate plus the output pixel rate. The *segs* part equals the horizontal resolution divided by 128 (rounded up).
- ▷ The quality equals the output frame rate.
- ▷ The parameters equal the video parameters of the output.

This results in 270 configurations for each Scaler component.

Every configuration of an application component has well-defined required budgets (determined during system design). The QRM aggregates from the provided physical platform budgets VEPs that provide budgets matching the required application budgets. The QRM also computes the bindings of the VEPs to the physical resources (Fiber/HW scaler compositions in this example).

Algorithm 1 shows the QRM algorithm for the video-processing system. This algorithm explores all configurations of the aggregation of stream and platform components, as illustrated in Fig. 2.1, and returns a cost-optimal configuration as a solution to the given multi-objective optimization problem. The algorithm is assumed to have access to all the interface models of the application and platform components (Transport (t), Scaler (sc), Fiber (f), HW Scaler (hw)). It has as parameters the number of execution-platform components k , n tuples (v_i, x_i) , where v_i is a *video* value that specifies the resolution and

Algorithm 1: *VideoProcessingQRM*($k, (v_1, x_1) \dots, (v_n, x_n), cost$)

```

1:  $R \subseteq \mathbb{N}^n := \emptyset$ 
2: for  $i = 1$  to  $k$  do
3:    $ex_i := f \parallel hw$ 
4: end for
5: for  $i = 1$  to  $n$  do
6:    $app_i := (t \Rightarrow sc) \cap D_{6, \{(v_i, x_i)\}}$ 
7:   create Connection  $c_i$ 
8:   create Virtual scaler  $vsc_i$ 
9:    $vep_i := c_i \parallel vsc_i$ 
10:   $va_i := vep_i \uparrow app_i$ 
11: end for
12: for all mappings  $m : \{1, \dots, n\} \rightarrow \{1, \dots, k\}$  do
13:   for  $i = 1$  to  $n$  do
14:     $ex_{m(i)} := ex_{m(i)} \uparrow va_i$ 
15:   end for
16:    $S := \pi \circ \min(\downarrow 2 \circ (\downarrow 1)^4 \circ (ex_1 \parallel \dots \parallel ex_k))$ 
17:    $R := \min(R \cup S)$ 
18: end for
19: return select one configuration from  $\min(cost(R))$ 

```

rate of the incoming stream i , and x_i is a *video params* value that specifies the required output resolution after scaling of stream i , and a *cost* function. The cost function is an \preceq -derivation on the configuration space \mathbb{N}^n that captures the output rate of each stream. The Pareto-minimal solutions are captured in set R , which is initialized on line 1. It contains the quality of each stream (the output rate in Hz). Lines 2–4 initialize the k execution-platform components. Lines 5–11 initialize the n application components and their VEPs. In line 6, we apply a subset constraint to select the proper parameters for the stream. The proper VEP for the application component is computed in lines 7–9, and in line 10 we create the vertical aggregation of the application and the VEP component. The for-loop in lines 12–18 iterates over all possible mappings of the n application components on the k execution-platform components. An entry m gives for an application component i the execution-platform component $m(i)$ to which it is bound via a VEP. Lines 13–15 then bind the application-VEP combinations to the execution-platform components using vertical aggregation. In each iteration of this loop, one application-VEP combination is bound to the (remaining) budget of the platform component on which it is mapped. After this loop, ex_i is an aggregation of an execution-platform component with zero or more application-VEP components. Line 16 combines all these ex_i aggregations through free aggregation. The abstractions filter out everything but the qualities of the streams. A permutation π is applied such that the quality of stream i is on the i -th position in the resulting configuration space. These configurations then are combined with the alternatives of other mappings in line 17. Finally, line 19 applies the cost function and selects one of the optimal solutions.

In principle, we enumerate all possible mappings of streams onto hardware components. Of course, this can be a huge number, and therefore we have applied *symmetry reduction* [CEJS98]. Since the execution-platform components are totally interchangeable, we can define an equivalence relation on the mappings m in line 12, and we only need to explore one

aggregation per equivalence class. For instance, if we have two hardware components, say ex_1 and ex_2 , and two streams as in line 10 of the algorithm, say va_1 and va_2 , then the mapping $ex_1 \uparrow va_1$ and $ex_2 \uparrow va_2$ gives the same result as the mapping $ex_2 \uparrow va_1$ and $ex_1 \uparrow va_2$ since ex_1 and ex_2 are interchangeable. We can also apply symmetry reduction to the streams by assuming that streams of the same type (determined by their input and required output) are completely interchangeable. Of course, the soundness of this second reduction depends on the cost function that is ultimately used. For instance, one stream may be more important for the user than another stream with equal parameters. This difference in priority is then typically expressed in the cost function, and clearly such streams are not interchangeable. Note that this application of symmetry reduction is an ad-hoc optimization specific for this case and the example QRM as shown in Alg. 1. A systematic approach to symmetry reduction is, for now, out of scope of our mathematical framework. In Sec. 8, we introduce a domain-specific language for the specification of the QRM interfaces of components and their composition. We can imagine that further developments of this language may support syntactic constructs that indicate symmetries in the search space.

Now consider the situation where $n = 6$, three HD+@90 to HD streams and three HD@60 to qHD streams, and $k = 3$. We use the cost function that maximizes the minimal rate. First, we observe the effect of symmetry reduction. Without symmetry reduction, we have 690 different mappings. Using symmetry reduction for the three execution-platform components results in 115 possibilities. Also using symmetry reduction for streams of the same *video* type yields only 18 possibilities. With or without stream symmetry reduction, we obtain a single configuration with the following quality at the end of the algorithm:

$$(90, 90, 90, 60, 60, 60)$$

This thus shows that we can process all six streams without rate loss. Each execution-platform component hosts an HD+@90 stream and an HD@60 stream. Now assume that we need to additionally scale a qHD@60 stream to nHD. The number of mappings with only symmetry reduction for the execution-platform components is 315, and the number of possibilities when also stream symmetry reduction is used, is 40. Optimization with both forms of symmetry reduction results in three Pareto-minimal configurations:

$$c_1 = (30, 90, 90, 60, 60, 60)$$

$$c_2 = (90, 90, 90, 20, 60, 60)$$

$$c_3 = (90, 90, 90, 15, 60, 60)$$

All three configurations come from the same mapping: two execution-platform components process an HD+@90 and an HD@60 stream each, and the third additionally processes the qHD@60 stream. The cost function now selects c_1 since it is the unique optimum configuration, i.e., it has a higher minimum output rate of 30 Hz among the streams than the other two configurations. Optimization without symmetry reduction for streams yields nine Pareto-minimal configurations at the end of line 18. Each of the three configurations above comes in three variations that correspond to permutations of the streams of the same type.

With seven streams in the system, and 270 configurations for each Scaler, the raw number of possible configuration combinations equals $270^7 = 104,603,532,030,000,000$. The total number of combinations of app_i configurations (line 6 of the algorithm; determined by the possible output rates for each stream) equals $6^7 = 279,936$. Clearly, a naive approach to optimization can end badly, that is, not end at all for all practical purposes. Using the

mathematical framework from this paper, our Java QRM implementation optimizes the seven streams using symmetry reduction in approximately 200 ms on a standard laptop computer (40 mapping possibilities). Without stream symmetry reduction, the average execution time is around 800 ms (315 possibilities). We believe that the implementation can still be improved significantly to reduce the running time. Our conclusion therefore is that our framework gives apt mathematical tools for runtime quality and resource management.

8. THE QRML SPECIFICATION LANGUAGE

In [vdBCH⁺20], we have introduced the Quality- and Resource-Management Language (QRML) and toolset. In this section, we illustrate (an improved version of) the component language for specification of QRM interfaces and their composition (online tooling available at <https://qrml.org>). Three fundamental parts of the mathematical framework presented in this paper are (i) partially-ordered sets to define the configuration spaces of the components, (ii) composition as an aggregation of alternatives, and (iii) optimization. In the next subsections, we explain how these are supported by the QRML toolset.

8.1. Partially-ordered sets in QRML. The language provides primitive sets of values (integers and reals) and enumerations, and supports the construction of combinations of these sets. Each set of values has a default partial order and the language provides means to manually specify a partial order (via lambda expressions). A QRML type is a set of values and its partial order, and these types are used to specify the configuration spaces of components. By convention, every type in QRML has a least and greatest element, denoted in the syntax by `bot` and `top` respectively.

```

budget Bw integer

budget FrameRate integer

budget Computation integer

channel Video {
  hres : integer ordered by =
  vres : integer ordered by =
  rate : integer ordered by =
} ordered by a<=b if a.hres <= b.hres & a.vres <= b.vres & a.rate <= b.rate

budget Scaling {
  segs : integer
  comp : Computation
} ordered by element-wise

budget Scalers {
  streams : integer
  scaling : Scaling
}

```

Figure 6: Type definitions for the video-processing system in QRML.

Some of the types for the video-processing example from Sec. 2 are shown in Fig. 6. A type definition starts with one of the keywords `typedef`, `channel` or `budget`. The `Bw`,

FrameRate and **Computation** budgets are specified as a partial order consisting of the integers and the (default) \leq relation. The **Video** channel type is a combination of three integers for the horizontal resolution, vertical resolution and frame rate, respectively. Note that the partial-order relations for each of these three parts is the equality relation on integers. This implies that, in this model, the values must match exactly and a higher provided resolution is not accepted to match a lower required resolution. The **ordered by** clause gives a predicate that specifies how two elements of the **Video** type are ordered by the partial-order relation. This definition uses the partial-order relations of the parts. Thus, $(h_1, v_1, r_1) \leq (h_2, v_2, r_2)$ if and only if $h_1 = h_2$ and $v_1 = v_2$ and $r_1 = r_2$. The **Scaling** budget also is a combination in which the **segs** part is ordered by the \leq relation on the integers, and the ordering of the **comp** part is given by the order on the **Computation** budget. The **element-wise** ordering is short-hand for the relation defined similarly as for the **Video** type (which can alternatively be defined with the **element-wise** clause). Finally, the **Scalers** budget shows another level of nesting of types. Because the order is not explicitly defined for the **Scalers** budget, its partial-order relation is the default element-wise comparison using the partial-order relations of the parts. Note that the inclusion of default orders in the semantics of QRML implies that the **ordered by** clauses for the **Video** and **Scaling** types could have been omitted. They are given here for illustration purposes.

8.2. Components and their composition in QRML. The types can be used to specify QRML components. A component declaration consists of three elements: port declarations, subcomponent declarations, and constraint declarations. Components have zero or more ports for each QRM part, and each port has a type. A port is declared by one of the keywords **provides**, **requires**, **input**, **output**, **quality**, or **parameter**. Multiple ports for the same part are combined in a multi-dimensional poset for that part using the free product. QRM parts with no explicit port specification default to the void poset. The port declarations of a component thus specify its configuration space. Components can have zero or more subcomponents. Subcomponents are declared by the **contains** keyword, and refer to components declared elsewhere. A **contains** statement can also offer alternatives, e.g., the component contains subcomponent *A* or subcomponent *B*. Finally, components can specify a number of constraints on the ports of the component and its subcomponents. These constraints are expressions on the port variables, and are specified between curly brackets after port declarations, via the **from** keyword after port declarations, or via the **constraint** keyword.

For instance, Fig. 7 shows how the execution platform of the example can be specified in QRML with the aggregation pattern. The **Fiber** and **HWscaler** components are atomic: they are not built from other components. They each have a single port and constraints on the values of their port are specified between curly brackets. The **ExecutionPlatform** component has two subcomponents declared by the **contains** keyword: a **Fiber** and a **HWscaler**. Furthermore, it declares two ports for the provided part of the QRM interface. The configuration spaces of the three components are the following:

$$\mathcal{S}_f = V \times V \times V \times (\mathbb{Z} \cup \{\perp, \top\}) \times V \times V$$

$$\mathcal{S}_h = V \times V \times V \times (\mathbb{Z} \cup \{\perp, \top\})^3 \times V \times V$$

$$\mathcal{S}_e = V \times V \times V \times (\mathbb{Z} \cup \{\perp, \top\})^4 \times V \times V$$

```

component Fiber {
  provides Bw p_bw { p_bw = 10*1000 }
}

component HWscaler {
  provides Scalers p_sc { streams = 4; scaling.comp = 300; scaling.segs = 32 }
}

component ExecutionPlatform {
  contains Fiber fiber
  contains HWscaler hwScaler

  provides Bw p_bw from fiber.p_bw
  provides Scalers p_sc from hwScaler.p_sc
}

```

Figure 7: Aggregation: the execution platform.

(where $V = \{\perp\}$ is the void poset that uses the $=$ partial order). Both the `Fiber` and the `HWscaler` have a single configuration which is specified by their respective port constraints. Their configuration sets are the following:

$$\begin{aligned}
 C_f &= \{(\perp, \perp, \perp, 10000, \perp, \perp)\} \subseteq \mathcal{S}_f \\
 C_h &= \{(\perp, \perp, \perp, (4, 300, 32), \perp, \perp)\} \subseteq \mathcal{S}_h
 \end{aligned}$$

The `ExecutionPlatform` specifies an aggregation. We use the constraint-based pattern for its semantics. The two `from` constraints relate the provided budget of the aggregate to the provided budgets of the constituents. The aggregation formally is the following instance of Eq. 5.4:

$$\min \circ (\downarrow 1)^{12} \circ \cap \{(x_1, \dots, x_{18}) \mid x_{16} = (x_4, x_{10})\} (C_f \times C_h \times \mathcal{S}_e)$$

Here, x_{16} is the provided budget of the `ExecutionPlatform`, x_4 is the provided budget of the `Fiber`, and x_{10} is the provided budget of the `HWscaler`. In this case, we thus obtain a single configuration for the `ExecutionPlatform` component with only a composite provided budget:

$$C_e = \{(\perp, \perp, \perp, (10000, 4, 300, 32), \perp, \perp)\} \subseteq \mathcal{S}_e$$

We use Ex. 6.1 to explain the alternatives construction in QRML. It introduces the `SWscaler` component that only provides the `Computation` budget. The `HWorSWscaler` component models the choice between the `HWscaler` and `SWscaler`. Fig. 8 shows the three relevant components. The configuration spaces for the `HWscaler`, `SWscaler` and `HWorSWscaler` follow from the port declarations and are the following:

$$\begin{aligned}
 \mathcal{S}_h &= V \times V \times V \times (\mathbb{Z} \cup \{\perp, \top\})^3 \times V \times V \\
 \mathcal{S}_s &= V \times V \times V \times (\mathbb{Z} \cup \{\perp, \top\}) \times V \times V \\
 \mathcal{S}_{hs} &= V \times V \times V \times (\mathbb{Z} \cup \{\perp, \top\})^3 \times V \times V
 \end{aligned}$$

The `contains or` construction in the `HWorSWscaler` means that this component provides a choice between the `HWscaler` and `SWscaler`, i.e., it follows the alternatives pattern. The configuration sets of the `HWscaler` and `SWscaler` are defined as in Ex. 6.1 and are denoted by C_h and C_s , respectively. Both consist of a single configuration. To compute the configurations of the `HWorSWscaler`, we take the alternatives of the two constituent components. The

```

component HWscaler {
  provides Scalers p_sc { streams = 4; scaling.comp = 300; scaling.segs = 32 }
}

component SWscaler {
  provides Computation p_cmp { p_cmp = 100 }
}

component HWorSWscaler {
  contains HWscaler hs or SWscaler ss

  provides Scalers p_sc

  constraint p_sc = hs.p_sc
  constraint p_sc = [top, ss.p_cmp, top]
}

```

Figure 8: Alternatives: a choice between components.

two constraints give the normalization constraints for Eq. 5.2. The alternatives operation formally has the following semantics:

- (1) $\min \circ (\downarrow 1)^6 \circ \cap \{(x_1, \dots, x_{12}) \mid x_{10} = x_4\} (C_h \times \mathcal{S}_{hs})$
- (2) $\min \circ (\downarrow 1)^6 \circ \cap \{(x_1, \dots, x_{12}) \mid x_{10} = (\top, x_4, \top)\} (C_s \times \mathcal{S}_{hs})$
- (3) $\min \circ \cup$

Here, x_{10} is the provided budget of the `HWorSWscaler`, and x_4 is the provided budget of the `HWscaler` in step 1 and the provided budget of the `SWscaler` in step 2. This results in two configurations for the `HWorSWscaler`:

$$S_{hs} = \{(\perp, \perp, \perp, (\top, 100, \top), \perp, \perp), (\perp, \perp, \perp, (4, 300, 32), \perp, \perp)\}$$

8.3. Optimization. We have defined the semantics of a QRML model in terms of configuration sets and Pareto-algebraic operations on these sets. Each component either is an atomic component (without subcomponents), a choice between subcomponents, or an aggregation of subcomponents. The first option directly gives the configuration set of the component through its configuration space and the constraints. The latter two options give the configuration set of the component through instances of the alternatives and aggregation patterns of Eq. 5.2 and Eq. 5.4, respectively. The Pareto-minimal configuration set of a component is sufficient and necessary for optimization with respect to an arbitrary cost function (i.e., an arbitrary \preceq -derivation). The actual computation of configuration sets of components is delegated to external tools through proper language transformations (this is work in progress). For instance, if all used constraints preserve dominance, then we can use Pareto-minimal configuration sets during all computations. This is a good match for a tool such as the Pareto calculator [GB07]. The Pareto-algebraic operations can be implemented sufficiently efficiently for on-line use. QRML specifications may then serve to configure a quality and resource manager, as the one introduced in the previous section. An alternative for design-time optimization, which does not require dominance preservation of operations, is a constraint-based formulation suitable for a solver such as Z3 [DMB08]. This formulation can then be used in an optimization approach such as presented in [LLGCM10] to approximate the Pareto-minimal set of configurations.

9. CONCLUSIONS

We have presented an interface-modeling framework for quality and resource management of component-based systems. The QRM interface of components is defined in terms of partially-ordered sets, and is founded on the Pareto Algebra of [GBTO07]. Our framework supports compositional refinement in the sense that a component interface C can safely be replaced with a component implementation C' if the latter dominates the former, i.e., if $C \preceq C'$. A running example of a video-processing system has shown how we can apply a number of basic operations (formally defined as functions) to create meaningful component compositions. The framework enables runtime quality and resource management in embedded and cyber-physical systems. Although we have defined several generally applicable composition patterns, typically domain or even application-specific knowledge is needed. This is certainly also true for the partially-ordered sets that form the foundation, and for the composition of the basic functions to model the composition of components. The framework that we have presented is therefore an abstract, conceptual framework. Concrete posets, constraints, derivations and ways of composition must be created for each application domain such that the semantics of that domain are reflected properly. We have shown how the QRML modeling language (<https://qrml.org>) can be used to specify such domain-specific quality and resource management models. QRML is under further development.

ACKNOWLEDGEMENTS

Twan Basten had the pleasure of being one of the first PhD students of Jos Baeten in Eindhoven. This experience had a lasting impact on his scientific work. Algebra is never far away. Twan and Jos have always shared an interest in describing and analyzing complex systems in mathematically rigorous and elegant ways. In a sense, the work in this paper found its inspiration already in Twan's PhD research. Throughout his career, Jos combined research on systems engineering with managing research organizations. Both involve balancing a variety of resources and qualities to achieve desired goals. An algebra for managing qualities and resources thus seems a fitting contribution to this *liber amicorum*.

REFERENCES

- [APA⁺16] A. Agirre, J. Parra, A. Armentia, E. Estévez, and M. Marcos. QoS Aware Middleware Support for Dynamically Reconfigurable Component Based IoT Applications. *International Journal of Distributed Sensor Networks*, 12(4), 2016.
- [BBS06] A. Basu, M. Bozga, and J. Sifakis. Modeling Heterogeneous Real-time Components in BIP. In *Proc. 4th IEEE International Conference on Software Engineering and Formal Methods, SEFM '06*. IEEE, 2006.
- [BCF⁺08] A. Benveniste, B. Caillaud, A. Ferrari, L. Mangeruca, R. Passerone, and C. Sofronis. Multiple Viewpoint Contract-Based Specification and Design. In *Formal Methods for Components and Objects, FMCO 2007*, LNCS 5382. Springer, 2008.
- [Box98] D. Box. *Essential COM*. Addison-Wesley, 1998.
- [CEJS98] E. M. Clarke, E. A. Emerson, S. Jha, and A. P. Sistla. Symmetry Reductions in Model Checking. In *Proc. Computer Aided Verification, CAV 1998*. Springer, 1998.
- [dAH01a] L. de Alfaro and T.A. Henzinger. Interface automata. *SIGSOFT Softw. Eng. Notes*, 26(5), 2001.
- [dAH01b] L. de Alfaro and T.A. Henzinger. Interface Theories for Component-Based Design. In *Proc. Embedded Software, EMSOFT 2001*, pages 148–165. Springer, 2001.

- [DMB08] L. De Moura and N. Bjørner. Z3: An Efficient SMT Solver. In *Proc. Tools and Algorithms for the Construction and Analysis of Systems*, TACAS 2008. Springer, 2008.
- [EJL⁺03] J. Eker, J. W. Janneck, E. A. Lee, J. Liu, X. Liu, J. Ludvig, S. Neuendorffer, S. Sachs, and Y. Xiong. Taming Heterogeneity - the Ptolemy Approach. *Proceedings of the IEEE*, 91(1), 2003.
- [EL07] S. A. Edwards and E. A. Lee. The Case for the Precision Timed (PRET) Machine. In *Proc. 44th ACM/IEEE Design Automation Conference*. ACM, 2007.
- [GB07] M. Geilen and T. Basten. A Calculator for Pareto Points. In *Proc. Design, Automation & Test in Europe Conference & Exhibition, DATE '07*. IEEE, 2007.
- [GBE10] M. García-Valls, P. Basanta-Val, and I. Estévez-Ayres. A Component Model for Homogeneous Implementation of Reconfigurable Service-based Distributed Real-time Applications. In *Proc. 10th Annual International Conference on New Technologies of Distributed Systems*, NOTERE '10. IEEE, 2010.
- [GBTO07] M. Geilen, T. Basten, B. Theelen, and R. Otten. An Algebra of Pareto Points. *Fundamenta Informaticae*, 78(1), 2007.
- [GKN⁺17] K. Goossens, M. Koedam, A. Nelson, S. Sinha, S. Goossens, Y. Li, G. Breaban, R. van Kampenhout, R. Tavakoli, J. Valencia, H. A. Balef, B. Akesson, S. Stuijk, M. Geilen, D. Goswami, and M. Nabi. NoC-Based Multiprocessor Architecture for Mixed-Time-Criticality Applications. In S. Ha and J. Teich, editors, *Handbook of Hardware/Software Codesign*. Springer, 2017.
- [GTW11] M. Geilen, S. Tripakis, and M. Wiggers. The Earlier the Better: A Theory of Timed Actor Interfaces. In *Proc. 14th Int. Conf. on Hybrid Systems: Computation and Control*, HSCC '11. ACM, 2011.
- [HM06] T. A. Henzinger and S. Matic. An Interface Algebra for Real-Time Components. In *Proc. 12th IEEE Real-Time and Embedded Technology and Applications Symposium*, RTAS '06. IEEE, 2006.
- [KLSV03] D.K. Kaynar, N.A. Lynch, R. Segala, and F. Vaandrager. Timed I/O Automata: A Mathematical Framework for Modeling and Analyzing Real-Time Systems. In *Proc. 24th IEEE International Real-Time Systems Symposium*, RTSS '03. IEEE, 2003.
- [KSHS17] I. Kurtev, M. Schuts, J. Hooman, and D.-J. Swagerman. Integrating Interface Modeling and Analysis in an Industrial Setting. In *Proc. 5th International Conference on Model-Driven Engineering and Software Development*, MODELSWARD '17. SciTePress, 2017.
- [LLGCM10] J. Legriél, C. Le Guernic, S. Cotton, and O. Maler. Approximating the Pareto Front of Multi-criteria Optimization Problems. In *Proc. Tools and Algorithms for the Construction and Analysis of Systems*, TACAS 2010. Springer, 2010.
- [LM87] E. A. Lee and D. G. Messerschmitt. Synchronous Data Flow. *Proceedings of the IEEE*, 75(9), 1987.
- [LMFS13] A. Luppold, B. Menhorn, H. Falk, and F. Slomka. A New Concept for System-level Design of Runtime Reconfigurable Real-time Systems. *SIGBED Review*, 10(4), 2013.
- [LT87] N.A. Lynch and M.R. Tuttle. Hierarchical Correctness Proofs for Distributed Algorithms. In *Proc. 6th Annual ACM Symposium on Principles of Distributed Computing*, PODC '87. ACM, 1987.
- [OA18] OSGi Alliance. OSGi release 7, 2018.
- [OMG12] Object Management Group. The Common Object Request Broker: Architecture and Specification Version 3.3, 2012.
- [SKB⁺15] S. Sinha, M. Koedam, G. Breaban, A. Nelson, A. Beyranvand Nejad, M. Geilen, and K. Goossens. Composable and Predictable Dynamic Loading for Time-critical Partitioned Systems on Multiprocessor Architectures. *Microprocessors and Microsystems*, 39(8), 2015.
- [SM97] Sun Microsystems. JavaBeans, 1997.
- [SVDP12] A. Sangiovanni-Vincentelli, W. Damm, and R. Passerone. Taming Dr. Frankenstein: Contract-Based Design for Cyber-Physical Systems. *European Journal of Control*, 18(3), 2012.
- [vdBCH⁺20] F. van den Berg, V. Camra, M. Hendriks, M. Geilen, P. Hnetyinka, F. Manteca, P. Sánchez, T. Bures, and T. Basten. QRML: A Component Language and Toolset for Quality and Resource Management. In *Forum on specification & Design Languages*, FDL 2020. IEEE, 2020.

- [vOvdLKM00] R. van Ommering, F. van der Linden, J. Kramer, and J. Magee. The Koala Component Model for Consumer Electronics Software. *Computer*, 33, 2000.
- [WT05] E. Wandeler and L. Thiele. Real-time Interfaces for Interface-based Design of Real-time Systems with Fixed Priority Scheduling. In *Proc. 5th ACM International Conference on Embedded Software*, EMSOFT '05. ACM, 2005.