

This is a postprint version of the following published document:

López Fernández, L., García, B., Gallego, M. & Gortázar F. (2017). Designing and evaluating the usability of an API for real-time multimedia services in the Internet. *Multimedia Tools and Applications*, 76(12), pp. 14247–14304.

DOI: [10.1007/s11042-016-3729-z](https://doi.org/10.1007/s11042-016-3729-z)

© 2016, Springer Science Business Media New York

Designing, implementing and evaluating an API for real-time multimedia services in the Internet

Luis López-Fernández¹ • Boni García¹ • Micael Gallego¹ • Francisco Gortázar¹

Abstract In the last few years, multimedia technologies in general, and Real-Time multimedia Communications (RTC) in particular, are becoming mainstream among WWW and smartphone developers, who have an increasing interest in richer media capabilities for creating their applications. The engineering literature proposing novel algorithms, protocols and architectures for managing and processing multimedia information is currently overwhelming. However, most of these results do not arrive to applications due to their complexity and to the lack of simple and seamless APIs enabling non-expert developers to leverage them. Interestingly, in this context in which APIs are the critical ingredient for reaching wide developer audiences, the scientific literature about multimedia APIs is scarce. In this paper we try to contribute to fill this lack by proposing the RTC Media API: a novel type of API making simple for developers the use of latest trends in RTC multimedia including WebRTC, Video Content Analysis or Augmented Reality. We provide a specification of such API and discuss how it satisfies a set of stringent requirements including programming-language agnosticism, adaptation to cloud environments, support to multisensory multimedia, etc. To conclude, we describe an implementation of such an API that has been created in the context of the Kurento open source software project, and present a study evaluating the API usability performed in a group of more than 40 professional developers distributed worldwide.

Keywords Media Server, Real-Time multimedia Communications, Application Programming Interfaces, WebRTC, Multimedia Processing, Multimedia Tools and Applications, Cognitive Dimensions of Notations.

1 Introduction

Analysts such as Marc Andreessen claim that “software is eating the world” stressing the importance of software-centered models into the economy and the transition of traditional business to software-based organizations [7]. This trend is permeating into all areas of IT (Information Technologies) including also multimedia industries. In the last few years, we have witnessed how multimedia technologies have been evolving toward software-centered paradigms embracing cloud concepts through different types of XaaS (Everything as a Service) models [16].

More recently, another turn of the screw is taking place thanks to the emergence and popularization of APIs (Application Programming Interfaces). This is perfectly summarized by Steven Willmott with his claim “software is eating the world and APIs are eating software” [75]. Software developers worldwide are getting used to create their applications as a composition of capabilities exposed through different APIs. These APIs are typically accessible through SDKs (Software Development Kits) and expose in an abstract way all kind of capabilities including device hardware, owned resources and remote third party infrastructures. This model, applied to cloud concepts, is quite convenient for individual developers and

□Luis López-Fernández
luis.lopez@urjc.es
Tel: +34 914 888 747
Camino del Molino S/N, 28943, Fuenlabrada, Spain

¹ Universidad Rey Juan Carlos, Spain

small companies, which have now the opportunity of competing with large market stakeholders without requiring huge effort investments and without needing to acquire hardware infrastructure or software licenses. Thanks to this, in the last few years, we are experiencing an explosion of innovation with thousands of new applications and services both for WWW and smartphone platforms that are being catalyzed by the rich and wide ecosystems of APIs made available to developers.

This trend towards the “APIfication” is also invading the multimedia arena and, very particularly, the RTC (Real-Time multimedia Communications) area. Initiatives such as WebRTC [41] are bringing audiovisual RTC in a standard and universal way to WWW users. The main difference between WebRTC and other popular video-conferencing applications is that WebRTC is not a service, but a set of APIs enabling WWW developers to create their customized applications using standard WWW development techniques.

WebRTC belongs to the HTML5 ecosystem and has awakened significant interest among the most important Internet and telecommunication companies. As opposed to other previous proprietary WWW multimedia technologies, it has been conceived to be open in a broad sense, both by being based on open standards and by providing open source software implementations. Currently, a huge standardization effort on WebRTC protocols is taking place at different IETF working groups (WGs), being the RTCWeb WG the most remarkable one [38]. In turn, WebRTC APIs are being defined and consolidated at the W3C WebRTC WG [76]. WebRTC standards are still under maturation stage and they might take some time to consolidate. In spite of this, most of the major browsers in the market already support WebRTC and it is currently available in billions of devices providing interoperable multimedia communications.

Hence, WebRTC is an opportunity for the creation of a next generation of disruptive and innovative multimedia services catalyzed worldwide through those emerging APIs. However, to reach this goal, the WebRTC ecosystem needs to evolve further. Basing on WebRTC browser capabilities, services can only provide peer-to-peer communications, which restrict use-cases to simple person-to-person calls involving few users. In order to enhance this model, server side infrastructures need to be involved. This is not new: as it is well known, the traditional WWW architecture is based on a three tier model [29] involving an application server layer and a service layer, this latter typically reserved to databases. In the same way, rich media applications also base on an equivalent three tier model where the service layer provides advanced media capabilities. The media component in charge of providing such capabilities is typically called media server in the jargon.

There is not a formal definition of what a media server is and different authors use the term with different meanings. In this paper, we understand that a media server is just the server side of a client-server architected media system. We concentrate our attention on RTC media servers, which are specialized in RTC media problems. Commonly, RTC media server capabilities consist on the following [47]:

- Group communication capabilities: These include mixing and forwarding. This type of media servers is called MCU (Multipoint Control Unit) [63] following the H.323 terminology and usually takes the form of Mixing Mixers or Selective Forwarding Units (SFU) [73]
- Media archiving capabilities: These are related to the recording of the audiovisual streams into structured or unstructured repositories and the ability to recover them later for visualization.
- Media bridging capabilities: This refers to attaining interoperability among networks or domains having incompatible media formats or protocols. Transcoders and IMS (IP Multimedia Subsystem) Gateways are among the most popular in this area.

Media servers are a critical ingredient for transforming WebRTC into the next wave of multimedia communications, and the availability of mature solutions exposing simple to use yet powerful APIs is a necessary requirement in that area. However, most standardization and implementation efforts are still concentrated at the client side, and server side technologies are still quite fragmented. Although there are a relevant number of WebRTC media servers available they do not provide coherent APIs compatible with WWW development models. Developing solutions with them typically requires expertise with low level protocols such as SIP [58], XMPP [59] or MGCP [5], on which average WWW developers do not have any experience. In addition to this, most state-of-the-art WebRTC media servers just provide the three basic capabilities specified above and are extremely hard to extend with further features. However, nowadays, many RTC services involve person-to-machine and machine-to-machine communication models and require richer multimedia processing capabilities such as computer vision, augmented reality, speech analysis and synthesis, etc.

In this paper we propose an evolution on current state-of-the-art RTC media servers presenting a new type of RTC API for media server control. This API addresses many limitations such as the ones described above and is aligned with WWW development principles, architectures and methodologies. The main contributions of this paper are threefold. First, we introduce the main concepts of the above mentioned API. Second, we present how developers may leverage it and create applications providing transparent interoperability among heterogeneous formats and protocols through a modular and extensible architecture. Third, we present an evaluation of the proposed API based on the Cognitive Dimensions of Notations (CDs) framework [32].

The remainder of this paper is as follows. Section 2 summarizes some approaches of RTC media servers and APIs available in the literature. Section 3 presents the proposed RTC Media API. Section 4 describes a survey in which our API is evaluated by means of a research questionnaire following the CD framework. The last section concludes this research with discussion, contributions of the study and suggestions for further work.

2 Related work

2.1 RTC media server control APIs

Media server technologies emerged in 90's catalyzed by the popularization of digital video services. Initial media servers were specialized into specific functions such as streaming [45], transcoding [65] and RTC for audio and video conferencing [8]. In this paper we concentrate on this latter category.

The popularization of video and audio conferencing made RTC media servers evolve through different types of standards, which include H.323 [67], where the media server role is played by elements such as the MCU (Multipoint Control Unit) and the IMS (IP Multimedia Subsystem), and where media servers are generically called MRF (Media Resource Function) [43]. These standards were conceived by operators and corporate communications solution vendors, who concentrated on the specificities of their infrastructures and not on the needs of developers. As a consequence, the involved media control interfaces were designed based on low level protocols and not on high level friendly APIs. Among such protocols we can find the IETF MGCP [12], which evolved later into the 3GPP H.248 [66] recommendation. These are based on binary formats, which are hard to understand, implement, debug and extend. Probably due to this, these protocols did not have much impact out of telecommunication providers.

More recently, the commoditization of RTC media server technologies brought increasing interest on more flexible mechanisms for media control. Several IETF WG emerged with the objective of democratizing them among common developers. As a result, further protocols such as MSCML [69], MSML [60] emerged providing the ability of controlling media server resources through technologies understandable and familiar to average developers such as XML [15].

Although these protocols are simpler to understand and integrate, developing application on top of them is still a cumbersome, complex and error prone process. Due to this, many stakeholders noticed that the natural tools used by developers are not protocols but APIs and SDKs. Hence, a number of initiatives emerged trying to transform the protocol-based development methodology into an API-based development experience providing seamless media server control through interfaces adapted to programming languages specificities and not to infrastructure characteristics. In particular, the Java platform was one of the first on integrating this philosophy by trying to reproduce the WWW development experience and methodology for the creation of RTC media enabled applications. A relevant activity in this area is JAIN (Java API for Integrated Networks), which issued several APIs for the signaling, control and orchestration of media capabilities. These include the JAIN SIP API [50], the JAIN SLEE [27] API and the JAIN MEGACO API [8]; this latter being specifically devoted to control media servers through the H.248 protocol. JAIN APIs did not permeated much out of operators, but their ideas inspired more popular developments such as the SIP Servlet APIs [44] for the signaling plane, and the Media Server Control API (aka JSR 309) [25] for the media plane, which have been more widely used for the development of RTC solutions for voice and video.

Among all these APIs, this paper is especially interested in the JSR 309. JSR 309 concepts were quite revolutionary at the moment because the API tried to fully abstract the low level media server control protocols and media format details. The objective was to enable developers to concentrate on application

logic. JSR 309 defined both a programming model and an object model for media server control through a northbound interface, but independent of media server control protocols and hence, without requiring any specific southbound protocol driver. JSR 309 does not make any kind of assumption in relation to the signaling protocol or to the call flow, which are left to the application logic.

From a developer's perspective, probably the most innovative concept of JSR 309 was the introduction of a mechanism for defining the media processing logic in terms of a topology. This mechanism is based on an interface called `Joinable`. In JSR 309, all objects having the ability to manipulate media (e.g. send, receive, process, archive, etc.) implement such interface, which has a `join` method enabling interconnecting such objects following arbitrary dynamic topologies. Hence, a specific media processing logic can be implemented by developers just joining the appropriate objects. As an example, if you want to create an application mixing two RTP (Real-time Transport Protocol) streams and recording the resulting composite into a file, you just need to join the appropriate objects with the appropriate topology. Taking into consideration that in JSR 309 the `NetworkConnection` is the class of objects capable of receiving RTP streams, that `MediaMixer` is the class of objects with mixing capability and that `MediaGroup` is the class with the ability of recording; the above mentioned media topology can be achieved just by joining two `NetworkConnection` instances to a `MediaMixer` instance, which in turn, is joined with a recording `MediaGroup`. This approach makes possible for developers to conceive their media processing logic as graphs of "black-box" *joinables*, which is a quite modular and intuitive mechanism for working in abstract terms with the complex concepts involved in RTC multimedia applications.

Another relevant innovation of JSR 309 is the introduction of media events. Thanks to this mechanism, the media processing logic held by a media server can fire events to applications through a publish/subscribe mechanism. This is very convenient for enabling applications to become media-aware meaning that complex processing algorithms at the media server can provide asynchronous information dealing with things happening inside the media, for instance DTMF (Dual-Tone Multi-Frequency) tones being detected, voice activity being present, and so on.

JSR 309 permeated into mainstream developer audiences as a suitable API for media server control following the typical three tier model [25]. However, in the last few years, the emergence of novel technologies and computation paradigms have made JSR 309 to show relevant limitations. For example, nowadays group videoconferencing services are evolving from Media Mixing models, which require relevant media processing, towards SFU (Selective Forwarding Unit) models, which are based on media routing [73]. JSR 309 is heavily adapted to Media Mixing and, due to this, most of its APIs assume that participants send/receive only one media stream to/from the media server. As a consequence, SFU models do not fit nicely into JSR 309 APIs. This is particularly a problem when all the streams of a group videoconference are multiplexed into a single RTP session, as happens typically on modern WebRTC SFU media servers supporting bundle RTP [40], because JSR 309 APIs do not provide any kind of mechanism for demultiplexing streams from a `NetworkConnection`. Moreover, in JSR 309 the API specification explicitly forbids several input `NetworkConnections` to be joined to a single output `NetworkConnection`, as an SFU router would require. Instead, they need to be joined first to a `MediaMixer`, which, in turn, can be joined to the output `NetworkConnection`.

When looking to other modern RTC technologies, we notice again that the JSR 309 design has limitations. For example, if we consider WebRTC W3C APIs [76], we may understand that they split endpoint capabilities into different functional blocks each of which is exposed through an abstract interface (e.g. `RtpSender`, `RtpReceiver`, `PeerConnection`, etc.) However, if we want to expose WebRTC media server capabilities through JSR 309 we need to accept that endpoints can only be represented through the `NetworkConnection` interface, which is extremely limited to support rich WebRTC capabilities such as `DataChannels` [10], Trickle ICE [39], simulcast [74], etc.

JSR 309 shows also drawbacks in relation to its extensibility. In JSR 309 it is possible to support new media object types using `MediaGroups`, however, configuration of these new types have to be done with Media Server specific descriptions as strings, which cannot be validated by the compiler. It is important to note that these new media object types cannot be `NetworkConnection`, only `MediaGroups`. This is a hard limitation because no other network protocol different than RTP (negotiated through SDP) can be incorporated. The ideal would be to allow supporting the creation of new object types in a similar way than the core types, with factory methods in `MediaSession` (e.g.

`createNetworkConnection`, `createMediaGroup`, etc.), but this is not possible as `MediaSession` is an interface defined in JSR 309 API and hence it cannot be modified by the API user.

Further limitations about JSR 309 include:

- A counter-intuitive asynchronous development model based on an obscure `joinInitiate` primitive, which is incompatible with modern Java mechanism for managing asynchrony such as futures, continuations or lambdas. This lack of clean asynchronous programming model makes JSR 309 difficult to adapt to reactive programming frameworks and languages that are very demanded by developers today such as Node.js or Scala.
- A complete lack of mechanisms for monitoring and gathering quality stats on media sessions. This is an essential ingredient for production systems.
- JSR 309 it is designed specifically for the Java language. It would be desirable a portable API that can be used in as more languages as possible.
- This API is specifically designed to control Media Servers for phone communications because it exposes concepts like Dialogs (Prompt and record, DTMF, VoiceXML dialog, etc.). For example, it is mandatory for an implementation to provide a player with the capability to detect audio signals in DTMF, but this kind of functionality is not very useful in web applications.

2.2 API evaluation and characterization

APIs are critical, non-optional and cross-cutting in the construction of modern software systems [35]. Programming is a hard mental work and developers need to deal with large amounts of information for writing satisfactory code. In that duty, APIs are the most critical ingredient, specially when dealing with distributed systems and enterprise frameworks. For example, recent works [37] show that API misuse is the single most prevalent cause of software defects.

Designing APIs consists on conceiving abstractions through types and interfaces so that they can be consumed seamlessly, efficiently and safely by application developers. This is quite a complex topic for which very little is known and which requires interdisciplinary knowledge combining cognitive psychology and software engineering. However, the responsibility of API design is typically assigned to development team members who often do not have expertise or training in this area and who, typically, are more concerned with implementation details than with usability.

In spite of the well-known importance of APIs, API design and evaluation has not been a mainstream research topic and only recently some light has been shed on this area. Early attempts to investigating APIs typically followed unstructured and ad-hoc approaches concentrating on the specificities of given technologies. For example, works have been published with guidelines and recommendations for API design in C# [72], Java [14] or C++ [56] and for ad-hoc evaluation of new programming languages [17].

Using another perspective, some authors concentrated on specific problems transversal to all APIs with independence on their underlying technologies. Some remarkable efforts on this area enabled to understand that, for instance, the factory pattern tends to generate usability problems [24] and that there is a systematic set of questions that developers have when learning new APIs [23]. All these efforts are relevant due to the talent of their authors to detect and isolate common patterns and practices, but they do not make possible to build a consistent and reusable methodology for the area.

During the last decade further authors have tried to systematize the problem of API design and usability evaluation from a holistic perspective. Different approaches have been created for this [2, 11, 20, 26]. However, in this area, the one which has gained highest popularity is the Cognitive Dimensions of Notations (CDs) framework [13, 33]. CDs is a framework for describing the usability of notational systems. In this context, a notational system typically consists of a collection of symbols made on some medium and which define a behavior (i.e. meaning) through some kind of structured interactions. Examples of notational systems include English text on paper, buttons on a WWW GUI or programing with API calls on an IDE. CDs allow designers of notational systems to evaluate their designs with respect to the impact they have on the users of those designs. For this, CDs provides a vocabulary of terms (i.e. dimensions) that can be used to characterize cognitive artifacts and which makes possible to establish comparisons and to discuss and investigate about the implications of design decisions on those artifacts. It is important to remark that the dimensions defined by the CDs framework are not associated

to any notion of goodness and their meaning for a particular usability objective needs to be evaluated in the context of the specific activity a user is engaged in.

In the context of API evaluation, the CDs framework is a powerful tool because it allows to compare users' expectations and designer's views of the APIs with what the system actually provides. For example, early users of the HTML notation probably expected to be able to modify their pages headings easily, whereas the language required one action per heading for doing so. This signaled an imperfection in HTML API's usability which probably drove to the introduction of CSS (Cascading Style Sheets). In CDs terms such resistance to changes is characterized through a dimension called *viscosity*. Hence, in this case, we would say that CSS decreased the viscosity of HTML APIs. A comprehensive description of CDs dimensions can be found in Blackwell et al [13].

The CDs framework has been criticized due to its theoretical and practical limitations. For example, Moody et al [48] claim that CDs do not provide a scientific basis because of several reasons:

- The dimensions are vaguely defined often leading to misinterpretation in applying them.
- The theoretical and empirical foundations of the dimensions are poorly defined.
- The dimensions lack clear operationalization (i.e. evaluation procedures and metrics), which mean they can be only applied in a subjective manner.
- It does not support evaluation, as the dimensions simply define properties of notations and are not meant to be either "good" or "bad".
- It does not support design: the dimensions are not design guidelines and issues of effectiveness are excluded from its scope.
- Its level of generality precludes specific predictions meaning that it is unfalsifiable and, hence, it cannot be considered to provide a scientific basis for evaluating anything.

In spite of these criticisms, most authors accept that, although CDs need further evolutions and improvements, they are today the most suitable tool for performing comparative API evaluation and that their methodological principles allow analyzing real-world development problems on controlled lab studies in quite an efficient and lightweight manner [33].

2.3 Contributions of this paper: the RTC Media API requirements

APIs are always designed for satisfying requirements that are implicitly or explicitly assumed by the designer. The creation of the API proposed in this paper, that we call not surprisingly the RTC Media API, was founded also on a set of commonly accepted implicit requirements [14] plus a number of explicit ones. Among the former we have simplicity, usability, security, self-documenting and consistency. The latter were identified in the course of several large research projects devoted to RTC media [49, 28] as essential needs that should be provided by any modern RTC media API but, as discussed above, are not available in any state-of-the-art technology. The creation of an API complying with these requirements is the main contribution of this paper. They include the following:

Seamless API extensibility through custom modules

We want developers to be able to plug additional capabilities to the API (e.g. processing algorithms, protocols, etc.) and to enable their consumption as if they were native API capabilities (i.e. without requiring different syntax or language constructs.) The mechanism we require for this is based on modules in the sense that every extension takes the form of a module artifact (e.g. a `.jar` file in the Java language, a `.js` file in JavaScript language, etc.) and that developers may plug the modules they wish at development time without requiring any further modification or configuration. Remark that, for the reasons specified in sections above, JSR 309 does not comply with this requirement.

Adaptation to WWW technologies and methodologies.

This requirement has two aspects. The first, and most important, is the need of our API to be adapted to novel RTC WWW technologies and very particularly to WebRTC [41]. The WebRTC architecture, based on heavy use of RTP bundle [40] and RTCP demultiplexing mechanisms [53] and requiring complex ICE [57] management techniques such as Trickle ICE [39] makes complex to comply with this requirement. Also as specified in sections above, JSR 309 is not compatible with this as the `NetworkConnection` is based on plain RTP. The second, is the need of the API to adapt to the typical WWW three tier development model [67]. This means that the RTC Media API should be usable for WWW developers with their common development, deployment and debugging techniques and tools. To some extent, this means that the RTC Media API should be perceived by WWW developers as any other of the APIs consumed in the application logic, such as database APIs or ESB (Enterprise Service Bus) APIs.

Full abstraction of media details (i.e. codecs and protocols)

Media representation and transport technologies are complex and require specialized knowledge that is not typically available for common developers. For maximizing productivity and minimizing development and debugging complexity the RTC Media API should hide all the low level details of such technologies through the appropriate abstractions. In doing so, these abstractions must maintain the appropriate expressiveness enabling the API semantics to provide to developers the ability of performing the required operations onto protocols and formats including payloading, depayloading, decoding, encoding, re-scaling, etc.

Programming language agnostic

In today's Internet, developers use a multiplicity of programming languages for creating their applications. In fact, the majority of applications are called "polyglot" because use different languages. The specific choice depends on factors such as the previous experience, the personal preferences, the tasks to be accomplished, the target platform or the required scalability. In this context, tying developers to a specific programming language may be perceived as inflexible and unfriendly. For this reason, the RTC Media API needs to be language agnostic and adapt to the most common programming languages used nowadays. Of course, the specific syntax of the API calls may differ depending on language specificities. However, this requirement indicates that, somehow, the constructs, basic mechanisms and programming experience needs to be the same across different languages. This means, for example, that a developer having the appropriate expertise for creating applications with a Java RTC Media API implementation should be able of doing so with a JavaScript implementation as long as the subtleties of the two languages are known.

RTC media topology agnostic

One of the main objectives of RTC Media Servers is to provide group communication capabilities to applications. Due to this, any useful RTC media API must consider this as a central aspect of its design by exposing the appropriate constructs for group communications. When looking to how RTC group communications are technically implemented, we can notice that they are based on a set of well-known RTP interconnecting topologies [73] among which the most common ones are Media Mixing Mixers (MMM), Media Switching Mixers (MSM) and Selective Forwarding Units (SFU). In short, MMMs are based on the principle of composing a single output media stream out of N input media streams, so that the final composite stream represents the addition of the N input streams. MMMs require decoding of the N input streams, the generation of the composite (e.g. linear adding in audio or matrix layout for video) and encoding to generate the output stream. Due to the performance cost of these operations MMM do not scale nicely. On the other hand, MSMs and SFUs do not perform any heavyweight processing and they just forward and route N incoming streams to M outgoing streams, reason why they have better scalability properties. Their only difference is that MSMs enable the N to M mapping to change dynamically while on SFUs it is static and the only possible operation is switching on/off forwarding on any of the output M streams.

Understanding the differences and appropriate usage scenarios of these topologies is complex and a source of extra complexity for application developers. Due to this, we include a requirement for our RTC Media API to transparently manage all the subtleties of this problem so that the most appropriate solution is provided transparently by the API. Remark that JSR 309 also tried to comply with this requirement through the "Joinable" mechanism making possible for developers to establish topologies just by joining sources with sinks. However, as explained above, both JSR 309, and equivalently JSR 79, are only compatible with MMM topologies and cannot manage the, by the way most popular, MSM or SFU models.

Advanced media QoS information gathering

QoS is critical in multimedia services. Some milliseconds of latency or jitter can be the difference between successful and unsuccessful applications [71]. For this reason, RTC media developers need to have the appropriate instrumentation mechanisms enabling seamless debugging, monitoring and optimization of applications. These requirements guarantees that our RTC Media API developers are able to access advanced QoS metrics of the streams including relevant information such as packet loss, bandwidth, latency or jitter. Remark that none of the above mentioned RTC media server APIs, including the JSR 309, provide this kind of capability.

Compatibility with advanced media processing capabilities

So far, most RTC media technologies and APIs have been concentrated on the problem of transport (i.e. taking media information on one place and moving it to other places.) This happened because the most prevalent use case for RTC is person-to-person communications, where end-users expect from technology to eliminate distance barriers (i.e. to maintain a conversation as if it were face-to-face.) However, during the last decade novel use cases involving person-to-machine and machine-to-machine communications are gaining popularity in different verticals such as video surveillance, smart cities, smart environments, etc. In all these verticals, going beyond plain transport is a relevant requirement. As an example, the number of low latency RTC video applications being used in security scenarios is skyrocketing. In all these applications the ability to integrate Video Content Analysis (VCA) capabilities through different types of computer vision algorithms is an unavoidable requirement [34]. In addition, modern media applications in areas such as gaming or entertainment complement VCA with another trending technology: Augmented Reality (AR), which is also having high demand from users [78]. As a result, we include our RTC Media API to provide full compatibility with these advanced processing techniques enabling their seamless integration and use.

Context awareness

In RTC media services, as in other types of services, context is becoming a relevant ingredient for providing added value to applications [1]. Context is somehow an ambiguous concept for which there is not yet a formal definition. However, most authors accept context as any kind of information that can be used for characterizing the situation of an entity [21]. The OMA (Open Mobile Alliance) has generated a formal definition of context through the NGSI standard [9] as a set of attributes that can be associated to an entity. When working with RTC media, the entity is most typically a RTC media session (e.g. a media call).

Considering this context definition, this requirement means that our RTC media API needs to be capable of consuming context for customizing and adapting end-user experience but, most important, it needs to be capable of extracting context attributes from the media communication itself. In other words, the part of the context dealing with the media itself (i.e. what the media content is and what it represents at any time) needs to be manageable by the proposed API.

Adapted to multisensory multimedia

Traditionally RTC media has referred to simple audiovisual streams comprising typically one video track and one or two (i.e. stereo) audio tracks. However, modern trends and technologies extend this to a new multisensory notion [52], where multisensory streams may comprise several audio and video tracks (e.g. Multi-view and 3D video) but may also enable the integration of additional sensor information beyond cameras and microphones (e.g. thermometers, accelerometers, etc.) Hence, we establish a requirement for our RTC Media API to be capable of managing such multisensory multimedia in as seamless and natural way.

Adaptation to cloud media servers

Cloud computing is permeating in all IT domains, including multimedia, as the de-facto standard for system deployment and management [79]. This trend is also permeating into the RTC media server arena [30], reason why we need to consider it in the definition of our API. Adapting the RTC Media API to cloud environments basically means to make it compatible with how a PaaS (Platform as a Service) media server works [70]. In other words, our API needs to be compatible with a new notion of distributed media server, which in opposition with traditional monolithic media servers, is distributed through a cloud environment and can elastically scale to adapt to end-users generated load.

3 Description of the proposed API: the RTC Media API

3.1 API specification

3.1.1 MediaObjects: MediaElements and MediaPipelines

Before providing a formal description of the RTC Media API, which is probably too harsh, let's introduce some simple initial concepts that might be helpful for the understanding of the basic mechanisms and philosophy behind our API. The RTC Media API is built on top of an object oriented model where the root of the inheritance hierarchy is the `MediaObject`. The `MediaObject` is only a holder providing

utility members (it is abstract and cannot be instantiated). The two main types inheriting from `MediaObject` are `MediaElement` and `MediaPipeline`.

The `MediaElement` is the main abstraction of the RTC Media API. Intuitively, a `MediaElement` can be seen as a black box implementing a specific media capability. In general, `MediaElements` receive media streams through sinks, send media streams through sources and, in the middle, do “something” with the media. There are two main subclasses of `MediaElements`: `Endpoints` and `Filters`. An `Endpoint` is always a `MediaElement` with the ability of communicating media with the external world. All media streams coming into an `Endpoint` sink are send out of the `MediaElement` through some kind of external interface (e.g. network interface, file system interface, etc.) In the same way, all media streams received from the external interface are published and made available to other `MediaElements` through the `Endpoint` source. `Filters`, on the other hand, do not communicate media streams with the external world. Their only function is to implement some kind of media processing. This can be simple transport (e.g. a pass-through filter) or may involve complex processing algorithms including computer vision or augmented reality.

`MediaElements` can be connected among each other by means of a `connect` primitive. When a `MediaElement` (let’s call it A) is connected to other `MediaElement` (say B), the media streams available at A’s source are feed to B’s sink. The connectivity of `MediaElements` works following quite intuitive and natural rules. First, a `MediaElement` source can be connected to as many `MediaElement` sinks as you want (i.e. a `MediaElement` can provide media to many `MediaElements`). Second, a `MediaElement` sink can only receive media from a connected source. Hence, connecting a source to a sink that is previously connected makes that sink to first disconnect from its previous source before being connected to the new one. Hence, application developers create their media processing logic just by connecting media elements following the desired topology.

Another interesting feature of `MediaElements` is that the `connect` primitive is overloaded to provide the ability of connecting just one track of those available on a media stream. The RTC Media API distinguishes three types of tracks: `AUDIO`, `VIDEO` and `DATA`. The two former correspond with the typical audio-visual component of a stream. The latter represents arbitrary sensor data whose semantics is application-dependent. The `DATA` component makes possible to integrate any kind of sensor data into media applications.

Just for illustration, some example of `MediaElements` follows:

- `RtpEndpoint`: it represents an `Endpoint` having the capability of sending and receiving media streams based on standards such as the RTP protocol [64], the AVP and AVPF RTP profiles [51, 62], and the SDP media session negotiation mechanisms [35].
- `WebRtcEndpoint`: it represents an `Endpoint` having the capability of sending and receiving WebRTC streams complying with the appropriate standards and drafts [4].
- `PlayerEndpoint`: it represents an `Endpoint` with the ability of reading streams from different sources, such as a file system, an HTTP resource or RTSP [61].
- `FaceOverlayFilter`: it consists of a `Filter` using the Haar [46] computer vision algorithm for detecting faces on a stream and overlying on top of them images with customized scales and offsets.

`MediaPipelines`, in turn, are just containers of `MediaElement` graphs. A `MediaPipeline` holds `MediaElements` that can connect among each other following an arbitrary and dynamic topology. `MediaElements` owned by one `MediaPipeline` cannot connect to `MediaElements` owned by another `MediaPipeline`. Hence, the `MediaPipeline` represents an isolated multimedia session from the perspective of the application.

To illustrate these concepts, let’s create a simple application. This application performs a full-duplex back-to-back call between two users and records their streams into a repository. Fig. 1 shows the corresponding pipeline.

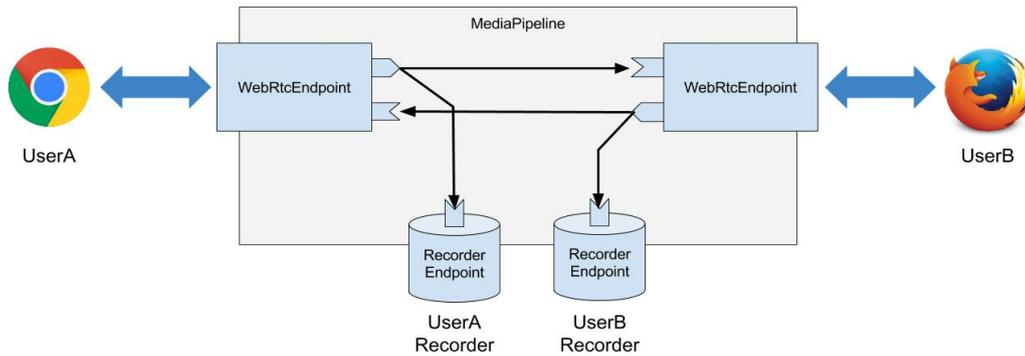


Fig. 1. Architectural diagram of an example application performing a back-to-back call between two users where their corresponding streams are recorded

This pipeline can be implemented in Java with the code shown in Table 1.

Table 1. Code snippet for developing the application specified in Fig. 1 in Java with the Kurento Client API, our reference implementation of the RTC Media API

```
KurentoClient rtcMediaAPI = KurentoClient.create();

MediaPipeline pipeline = rtcMediaAPI.createMediaPipeline();

WebRtcEndpoint userA = new WebRtcEndpoint.Builder(pipeline).build();
WebRtcEndpoint userB = new WebRtcEndpoint.Builder(pipeline).build();

userA.connect(userB);
userB.connect(userA);

RecorderEndpoint userARecorder =
    new RecorderEndpoint.Builder(pipeline, "videoUserA.webm").build();

RecorderEndpoint userBRecorder =
    new RecorderEndpoint.Builder(pipeline, "videoUserB.webm").build();

userA.connect(userARecorder);
userB.connect(userBRecorder);
```

3.1.2 RTC Media API IDL specification

One of the main requirements of the RTC Media API is that it should be available in different programming languages. Due to this, RTC Media API capabilities are specified through an IDL (Interface Definition Language) which is language agnostic. From an implementation perspective that IDL is compiled later to different programming languages in order to generate the appropriate SDKs. In this way, RCT media API capabilities are defined only once but the corresponding implementations can be generated for a variety of languages.

For simplicity, we have decided the RTC Media API IDL to be based on a JSON notation. In an RTC Media API file there are four sections: `remoteClasses`, `complexTypes`, `events` and `code`:

- The `remoteClasses` section is used to define the interface to media server objects. We call it “remote” because these objects are remote from the perspective of the API consumer, as they are hosted into the RTC media server. For example, `PlayerEndpoint` and `ImageOverlayFilter` are defined in this section in their corresponding IDL file.
- The `complexTypes` section is used to define enumerated types and registers used by remote classes or events. For example, the enumerated type `MediaType` with possible values `AUDIO`, `DATA` or `VIDEO` may be defined in this section.
- The `events` section is used to define the events that can be fired when using RTC Media API. For example, `EndOfStream` may be defined in the `events` section of the IDL file describing a `PlayerEndpoint`, so that the event is fired when the end of the stream is reached by the player.
- The `code` section is used to define properties to control the code generation phase for different programming languages. For example, in this section we can specify the package name in which all artifacts are generated for the Java language.

The code snippet shown in Table 2 outlines an example of an IDL file. For the sake of simplicity, we have replaced with dots (...) some parts of it.

Table 2. Example of an RTC Media API IDL file defining a `PlayerEndpoint` media element capability

```
{
  "code": {
    "api": {
      "java": {
        "packageName": "org.kurento.client",
        ...
      }
    }
  },
  "remoteClasses": [
    {
      "name": "PlayerEndpoint",
      "extends": "UriEndpoint",
      "constructor": {
        "params": [
          { "name": "mediaPipeline", "type": "MediaPipeline" },
          { "name": "uri", "type": "String" }
        ]
      },
      "properties": [
        { "name": "position", "type": "int64" }
      ],
      "methods": [
        { "name": "play", "params": [] }
      ],
      "events": [ "EndOfStream" ]
    },
    ...
  ],
  "events": [
    {
      "name": "EndOfStream",
      "extends": "Media",
      "properties": []
    },
    ...
  ],
  "complexTypes": [
    {
      "name": "MediaType",
      "typeFormat": "ENUM",
      "values": [ "AUDIO", "DATA", "VIDEO" ]
    },
    {
      "name": "Fraction",
      "typeFormat": "REGISTER",
      "properties": [
        { "name": "numerator", "type": "int" },
        { "name": "denominator", "type": "int" }
      ]
    },
    ...
  ],
  ...
}
```

As it can be observed, to define a remote class in Media API IDL it is mandatory to assign it a name. In addition, the following fields can be incorporated:

- **Extends:** A remote class may extend another remote class. In this case, all properties, methods and events of the superclass are available in objects of the subclass. Note that constructors of the superclass are not inherited. That is, they cannot be used to create objects of the subclass.
- **Constructor:** A remote class constructor is defined with a parameter list. Every parameter has a name and a type. The available types are: primitive types (`String`, `boolean`, `float`, `double`, `int` and `int64`), remote classes or complex types. Parameters can be defined as optional.
- **Properties:** A property is a value associated with a name. To define a remote class property it is necessary to specify its name and type. Properties can be defined as “read only”.

- **Methods:** Methods are named procedures that can be invoked with or without parameters. Every parameter is specified by its name and type. Parameters can be defined as optional. A return type can be specified if the method returns a value.
- **Events:** If a remote class declares an event it means that events of this type can be fired by objects of this remote class. It depends on the target programming language how this events are processed.

Remote classes are used mainly to define the `MediaElements` of the RTC Media API. To define a new `MediaElement` the only requirement is to define a new remote class that extends the built-in `MediaElement` remote class. This super class define the properties, methods and events of all `MediaElements`. The `MediaElement` class extends the `MediaObject` class, creating the class hierarchy represented in Fig. 2

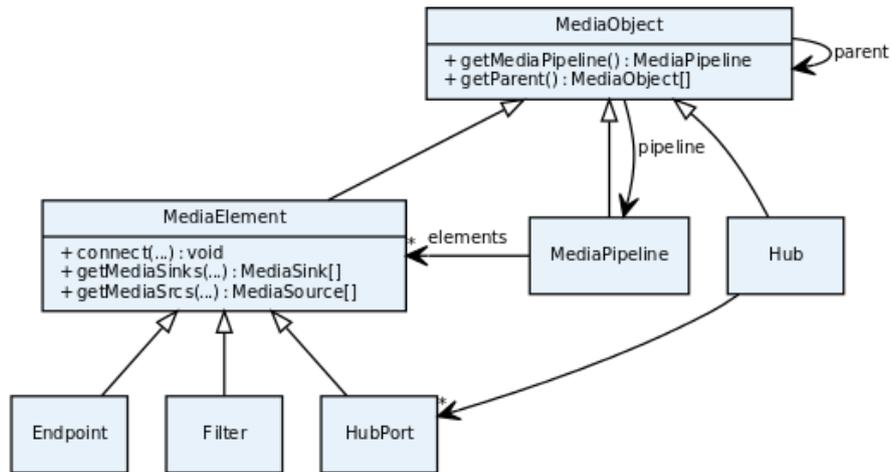


Fig. 2. `MediaObject` UML (Unified Modeling Language) inheritance diagram as defined in the RTC Media API IDL specification

To define an event, it is mandatory to assign it a name. In addition, an event can have properties. Every property must be defined with a name and a type. In the same way than remote classes, events can also extend a parent event type inheriting all its properties.

Regarding complex types, they can have two formats: enumerated or register. If a property or param is defined with an enumerated complex type, it can only hold a value from the list of specified values. For example, properties based on the enumerated complex type `MediaType` of Table 2 must have the value `AUDIO`, `DATA` or `VIDEO`. On the other hand, register complex types can hold objects with several properties. For example, the register complex type `Fraction` has two `int` properties: `numerator` and `denominator`.

To conclude, the code section is used to specify language-dependent configurations to the IDL compiler. Every programming language has its own section to avoid collisions. For example, the Java package name of the generated code has only sense in Java, while the name of the node module has only sense in JavaScript.

3.1.3 Compiling the RTC Media API IDL.

The IDL format described above makes possible to define the RTC Media API modules in a language-agnostic way. However, this needs to be translated into programming-language-dependent interfaces in order to have the real APIs to be used by application developers. The IDL compiler performs that task. Hence, we need to specify how this compilation happens so that all compiler implementations maintain compatibility on the generated code. For illustration, we have created such specification and as well as the compilers for the two most popular programming languages in the WWW: Java and JavaScript.

The Java IDL compiler works in the following way:

- **Package:** all artifacts (i.e. classes, interfaces and enums) are generated in the package specified in `code.api.java.packageName` section of JSON IDL file.
- **Remote classes:** For every remote class there are two generated artifacts: an interface and a builder class:

- Interface: For every remote class a Java interface is generated. This interface has the remote class methods defined in the IDL. In addition, for every property, a getter method is also included. The name of the method is the string “get” followed by property name. If the property is not read only, a setter method is also generated following the same approach. Finally, for every event declared in the remote class, a method to subscribe listeners to it is generated. For example, the `PlayerEndpoint` has the event `EndOfStream` declared in the IDL so the method `String addEndOfStreamListener(Listener<EndOfStream> listener)` is generated. The complementary method to remove the subscription is also generated. `Listener<E>` is a generic interface with only one method: `onEvent(E event)`.
- Builder class: We use the builder pattern [30] to create new remote class instances. A Builder is generated for each remote class. All mandatory params in the remote class constructor are mapped to parameters to the only constructor of the builder class. In this way, the compiler enforces that all mandatory parameters have a value. Optional constructor parameters are generated in builder class as fluent setter methods (prefixed with “with” instead of “set” or not prefixed if the method starts with “use”). The builder class is generated as an internal type of the above-mentioned interface to associate easily the class and the interface. The code snippet on Table 3 shows the creation of a `PlayerEndpoint` with the optional constructor parameter `useEncodedMedia` set to true.

Table 3. Code snippet showing how to instantiate a `PlayerEndpoint` in Java

```
PlayerEndpoint player =
    new PlayerEndpoint.Builder(pipeline, "video.webm").useEncodedMedia().build();
```

- Complex types: Depending on the complex type format (enum or register) the code generation is different:
 - Enumerated complex type: A Java enum class is generated.
 - Register complex type: A basic Java bean class is created. For every property, a getter and setter method is generated. In addition, a constructor with all properties as parameters is also generated. The code snippet in Table 4 shows a sample code using a register (`WindowParam`) as a constructor parameter of a `PointerDetectorFilter` remote class.

Table 4. Example illustrating how to instantiate a register complex type (`WindowParam`) as a Java bean

```
PointerDetectorFilter pointer = new PointerDetectorFilter.Builder(pipeline,
    new WindowParam(5, 5, 30, 30)).build();
```

- Events: For each event defined in a RTC Media API IDL file a new Java class is generated. “Event” is appended to the name of the class. This class is very similar to the generated classes for register complex types. That is, a getter and a setter method is included for each property. In addition, all event classes extend from the `RaiseBaseEvent` base class. This base class contains properties for holding the source of the event (`source`) and the timestamp in which the event was generated (`timestamp`). The code snippet in Table 5 shows an example illustrating how to work with events.

Table 5. Example illustrating how to work with events both in Java 7 and Java 8

```
PlayerEndpoint player =
    new PlayerEndpoint.Builder(pipeline, "video.webm").useEncodedMedia().build();

//Java 7
player.addEndOfStreamListener(new EventListener<EndOfStreamEvent>() {
    public void onEvent(EndOfStreamEvent e) {
        System.out.print("EOS Player "+e.getSource().getId());
    }
});

//Java 8
player.addEndOfStreamListener(
    e -> System.out.print("EOS Player "+e.getSource().getId())
);
```

When working with JavaScript IDL compilers, equivalent rules have been created:

- Package: We base on the NPM (Node Package Manager) [68] JavaScript packaging system. NPM mandates that a `package.json` file is generated. The following values are used:
 - o package name: `code.api.js.nodeName`
 - o package description: `code.api.js.npmDescription`
- Remote classes: For every class a new JavaScript prototype based class is generated. This class has all methods defined in the IDL file. In addition, for every property, a getter method is generated. Also setter methods are generated for non-read only properties. All generated methods have the parameters defined in the IDL plus a callback function. That callback parameter is used to implement the asynchronous execution of the method given that the API primitive may require to communicate with the RTC media server and, hence, cannot be synchronous. To create an object from a remote class, a factory method called `create` available at the pipeline object needs to be executed. The first parameter of the method is the name of the remote class to create as a string. The second is an options bag used when constructor parameters are required. The third, and last, is the async callback to receive the new object handler or an error. The code snippet in Table 6 shows the creation of a `PlayerEndpoint` with the mandatory parameter `uri` and optional constructor parameter `useEncodedMedia` set to `true`. As it can be observed, media element creation is an async operation.

Table 6. Code snippet showing how to instantiate a `PlayerEndpoint` in JavaScript

```
pipeline.create("PlayerEndpoint",{uri:"video.webm",useEncodedMedia:true},
  function(error, player) {
    if (error) return console.error(error);

    //use player here
  });
```

- Complex types: For enumerated complex types, there is no code generation. Enum values are simply strings. On the other hand, register complex type are generated as JavaScript classes based on prototypes. Also, for every registered complex type a factory function is generated to allow the creation of objects. The code snippet in Table 7 shows the creation of a `PointerDetectorFilter` using a complex type `WindowParam` as parameter.

Table 7. Example illustrating how to instantiate a register complex type (`WindowParam`) as a JavaScript object

```
var options = {
  "windowParam":WindowParam({
    "topRightCornerX":5,
    "topRightCornerY":5,
    "height":30,
    "width":30
  })
};
pipeline.create("PointerDetectorFilter",options,
  function(error, pointer) {
    if (error) return console.error(error);
  });
```

- Events: There is no classes generated for events in JavaScript. When an event is raised, a new object is created and populated with all relevant information as properties. In Table 8 a `Player` is created and a listener is registered for its event `EndOfStream`. When this event is generated, a function is executed with the event as parameter. This event parameter can be used to obtain the relevant information such as timestamp, source of the event, etc.

Table 8. Example illustrating how to work with events in JavaScript

```
pipeline.create("PlayerEndpoint",{uri:"video.webm",useEncodedMedia:true},
  function(error, player) {
    if (error) return console.error(error);

    player.on("EndOfStream", function(e){
      console.log("EOS player "+e.source.id);
    });
  });
```

3.1.4 Creation and deletion of media capabilities

Java and JavaScript have notable differences in media object creation. This is due to the differences in the type safety of both languages. Java is strongly typed. Hence, it is important that the compiler enforces typing in several contexts: mandatory parameters, optional parameters, media object signature, etc. On the other hand, in JavaScript there is no type checking until runtime and this is why we do not enforce any kind of protection.

The releasing of media objects is simple. We consider that a media object is released when the `release` method is invoked. In Java, the `release` method can be executed in a synchronous way, blocking the invoking thread until a response is received. That response can be successful or fail. In the latter case, an exception is thrown. In JavaScript, it is executed asynchronously. For this reason, a callback parameter is necessary so that failures can be notified. The following piece of code shows the release of a media object in Java and in JavaScript.

Table 9. Code snippets showing how to release a `PlayerEndpoint` both in Java and JavaScript

```
//Media object release in JavaScript
player.release(function(error) {
    if (error) return console.error(error);
});

//Media object release in Java
player.release();
```

3.1.5 Synchronous and asynchronous programming models in the RTC Media API

One of the most critical design decisions when designing APIs is how they behave in relation to threads. When performing I/O (Input/output) operations, there is a common agreement that asynchronous APIs are more scalable than synchronous ones [6]. Synchronous I/O typically block threads until a response is received or a timeout is reached. Hence, given that there is a practical limit on the number of threads in a system (mainly due to memory constraints), synchronous API models tend to generate thread starvation and decrease performance due to the overload they generate into the operating system task scheduler. To solve this problem, many modern APIs provide asynchronous I/O operations. In this case, the thread executing the I/O is not blocked after the invocation and can be used to execute other tasks. However, asynchronous APIs are more complex to use and are susceptible of suffering a problem called “callback hell” [42]. This is a well-known problem that arises when asynchronous calls are invoked in the callbacks of another asynchronous calls, creating a deep nesting of callbacks.

When we designed the RTC Media API we decided to provide developers the flexibility of choosing between the synchronous and the asynchronous models so that they were not limited by any of their corresponding drawbacks. Due to this decision, our Java IDL compiler generates two methods for each I/O operation: the synchronous and the asynchronous versions. Synchronous methods block the calling thread until a response is received. This can be appreciated, for example, in the code snippet shown in Table 6. After that, the execution continues. The asynchronous primitives, in turn, include as a last parameter a continuation, that is, an object that have two methods: `onSuccess`, that is executed when the response is received, and `onError`, that is executed when an error or timeout occurs. The code snippet in Table 10 shows an example.

Table 10. Example illustrating the creation of a `PlayerEndpoint` using the asynchronous Java API

```
new PlayerEndpoint.Builder(pipeline, "video.webm")
    .buildAsync(new Continuation<PlayerEndpoint>() {

        @Override
        public void onSuccess(PlayerEndpoint player) throws Exception {

            player.play(new Continuation<Void>() {

                @Override
                public void onSuccess(Void result) throws Exception {
                    log.info("Play started");
                }

                @Override
                public void onError(Throwable error) throws Exception {
```

```

        log.error("Exception invoking play in player", error);
    }
    });
}

@Override
public void onError(Throwable error) throws Exception {
    log.error("Exception creating player", error);
}
});

```

When going to JavaScript, things are more complex. Due to the characteristics of the JavaScript language both in the browser and in Node.js [68], only asynchronous I/O operations are possible. Due to this, and as it can be seen in the code snippets shown in Table 11, our IDL compiler includes a callback as the last parameter that is executed asynchronously when the operation is resolved. However, providing only this mechanism reduces the flexibility of developers to avoid the callback hell. This is why we designed novel mechanisms for simplifying developers' work. The first one is based on Promises [42]. A Promise represents an operation that has not completed yet, but is expected to do so in the future. Hence, an asynchronous method can return a promise object instead of expect a callback as last parameter. The developer specifies the code to be executed when the promise is fulfilled, executing a method called "then" with the callback as parameter. Table 11 shows a JavaScript code creating a player and invoking the play method on it, comparing the traditional implementation based on callbacks with an implementation using promises.

Table 11. Different forms of PlayerEndpoint creation in JavaScript (with callbacks, with promises, and with ES6 arrow functions)

```

//Asynchronous API with callbacks
pipeline.create("PlayerEndpoint",{uri:"video.webm",useEncodedMedia:true},
    function(error, player) {
        if (error) return console.error(error);

        player.play(function(error){
            if (error) return console.error(error);
            console.log("Play started");
        });
    });

//Asynchronous API with promises
pipeline.create("PlayerEndpoint",{uri:"video.webm",useEncodedMedia:true})
    .then(function(player){ return player.play(); })
    .then(function(){ console.log("Play started"); })
    .catch(function(err){ console.error(error); });

//Asynchronous API with promises and ES6 arrow functions
pipeline.create("PlayerEndpoint",{uri:"video.webm",useEncodedMedia:true})
    .then(player => player.play())
    .then(() => console.log("Play started"))
    .catch(err => console.error(error));

```

Moreover, if promises are combined with generators [55], a new ES6 (ECMAScript 6) feature, the asynchronous code can look like synchronous one. For illustration, remark that Table 12 code snippet implements the same logic than the one on Table 11, but using generators. As it can be observed, the improvement on code readability is noticeable.

Table 12. Creation of a PlayerEndpoint using generators in JavaScript. As it can be observed, the code readability improves significantly and the callback hell is fully avoided

```

co(function* () {
    try{
        var player = yield pipeline.create('PlayerEndpoint');
        yield player.play()
        console.log("Play started");
    } catch(e) {
        onError(e);
    }
}) ();

```

The next version of JavaScript, ES7, which is still under standardization, has a proposal to simplify this: the `async/await` keyword, which marks when a call is asynchronous but accepts synchronous API syntax. Using it, the code in Table 12 can be written as shown in Table 13 with ES7. As it can be observed, the `yield` keyword is replaced by `await` and the `co` function is no longer necessary.

Table 13. Creation of a `PlayerEndpoint` in ES7

```

try{
  var player = await pipeline.create('PlayerEndpoint');
  await player.play()
  console.log("Play started");
} catch(e) {
  onError(e);
}

```

3.1.6 RTC Media API capabilities

Once we have presented the formal aspects of the RTC Media API, we can switch to a more practical perspective and introduce its media capabilities. These capabilities comprise specific media objects that are made available to application developers to create their RTC media enabled applications following the above-described API guidelines. These capabilities can be grouped into two main categories: media elements, which inherit from the `MediaElement` class and manage a single media stream, and hubs, which inherit from the `Hub` class and have been specifically designed for the management of groups of streams.

As specified in sections above, media elements have two flavors: Endpoints and Filters. Endpoints are in charge of the I/O media operations in the media pipeline. Fig. 3 shows the RTC Media API endpoint inheritance hierarchy, which comprises the following capabilities:

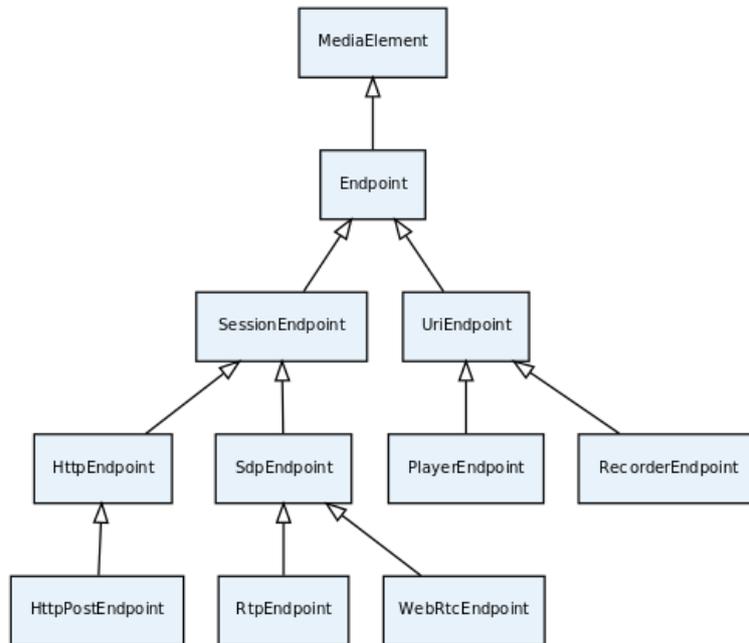


Fig. 3. UML class diagram of the Endpoints specified by the RTC Media API

- The `WebRtcEndpoint` is an I/O endpoint that provides full-duplex WebRTC media communications compatible with the corresponding protocol standards [4]. It is important to remark, that among `WebRtcEndpoint` capabilities, the RTC Media API defines as mandatory the `DataChannel` support. `DataChannels` are a mechanism for receiving media information beyond audio and video given their ability to accommodate arbitrary sensor data that is transported in the same ICE connection than the audio and the video and, hence, may maintain synchronization with them.
- The `RtpEndpoint` is equivalent but with the plain RTP protocol.
- The `HttpPostEndpoint` is an input-only endpoint that accepts media using HTTP POST requests. This capability needs to support HTTP multipart and chunked encodings, so that it is

compatible with the HTTP file upload function exposed by WWW browsers. This endpoint must support the MP4 and WebM media formats.

- The `PlayerEndpoint` is an input-only endpoint that retrieves content from the local file system, HTTP URLs or RTSP URLs and injects it into the media pipeline. This endpoint must support the MP4 and WebM media formats for all input mechanisms as well as RTP/AVP/H.264 for RTSP streams.
- The `RecorderEndpoint` is an output-only endpoint that provides function to store contents in reliable mode (doesn't discard data). This endpoint may write media streams to the local file system, or to HTTP URLs using POST messages. This endpoint must support MP4 and WebM media formats.

Filters, in turn, are used for processing media streams. Filters are useful for integrating different types of capabilities such as Video Content Analysis (VCA), Augmented Reality (AR) or custom media adaptation mechanisms. The RTC Media API does not specify any kind of mandatory filter and it is let to API implementers to define their filters following the RTC Media API extensibility mechanisms.

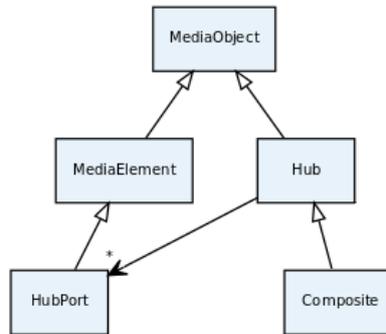


Fig. 4. UML class diagram of main Hub types in the RTC Media API

To conclude, hubs follow the inheritance scheme depicted in Fig. 4. Hubs work in coordination with `HubPorts`: a special type of media element, which provides sinks and sources to hubs. The RTC Media API only defines as mandatory hub type the `Composite`, which implements a MMM media topology, as described in previous sections. Developing with `Composites` is simple as long as the following rules are taken into account.

- `Composites`, as all hubs, act as a factory of `HubPorts`. This means that at a `Composite` instance we can create as many `HubPorts` as we want. These `HubPorts` are media elements having sources and sinks, which makes possible to connect other media elements to them and get media into and out of the hub.
- A `Composite` mixes all streams received from its `HubPort`'s sinks and exposes the resulting mixed stream at the sources. The audio of the mixed stream obtained at a `HubPort`'s source includes all the inputs except the one of its own `HubPort`'s sink. The video, on the other hand, combines all `HubPort`'s sinks into the resulting composite matrix.

3.1.7 Extending the API

One of the main requirements of the RTC Media API is extensibility: developer should be able to include new `MediaElements` into the API so that they maintain compatibility with other `MediaElements` defined natively by the API or by third parties. In order to support extensibility, we have created the notion of the RTC Media Module. A module is a bundle composed by:

- A module definition: the `MediaElements` interfaces and related types defined in the RTC Media API IDL.
- The corresponding software libraries: the specific language-dependent SDK enabling developers to use the module in their software projects.

According to this, imagine that you have created a new capability in your RTC media server involving some kind of computer vision algorithm to process a video stream and mark some relevant regions on it. The details about how this is implemented are out of the scope of this paper. The point is that, for exposing this new feature through the RTC Media API, the best choice is to create a `Filter`. Without loss of generality imagine we call it `CompuVisionFilter`. This filter interface needs to be defined in a RTC Media Module Definition file, which contains the RTC Media API IDL. If we suppose that the

filter requires an `int` parameter upon construction for tuning the behaviour of the algorithm and that it has a method that can be invoked at any time to enable or disable the processing (the “enable” method); the resulting module definition is the one shown in Table 14.

Table 14. Example of module definition for a new Filter called `CompuVisionFilter`

```
{
  "code": {
    "api": {
      "java": {
        "packageName": "com.company.filter"
      }
    }
  },
  "remoteClasses": [
    {
      "name": "CompuVisionFilter",
      "extends": "Filter",
      "constructor": {
        "params": [
          { "name": "mediaPipeline", "type": "MediaPipeline" },
          { "name": "param", "type": "int" }
        ]
      },
      "methods": [
        {
          "name": "enable",
          "params": [{ "name": "enabled", "type": "boolean" }]
        }
      ]
    }
  ]
}
```

From this file, the Java IDL compiler should be able to generate the `CompuVisionFilter` SDK library. When generated, the filter can be incorporated into any pipeline and interoperate with any of the rest of RTC Media API capabilities. Table 15 shows a code snippet illustrating how to create an application processing a media clip with the filter and exposing the resulting stream in real-time to a WebRTC capable browser. Remark that the example requires the filter to interoperate with built-in RTC Media API capabilities such as the `PlayerEndpoint` or the `WebRtcEndpoint`.

Table 15. Using the filter `CompuVisionFilter` previously defined

```
import com.company.CompuVisionFilter;

...

MediaPipeline p = rtcMediaApi.createMediaPipeline();

WebRtcEndpoint webrtc = new WebRtcEndpoint.Builder(p).build();
PlayerEndpoint player = new PlayerEndpoint.Builder(p,"file://...").build();

CompuVisionFilter filter = new CompuVisionFilter.Builder(p,10).build();

player.connect(filter);
filter.connect(webrtc);

//At any time...
filter.enable(true);
```

3.1.8 Implementing the RTC Media API: the Kurento Client API

In order to implement the RTC Media API and making it expose useful capabilities to developers we just need two ingredients. The first is an RTC media server. This media server needs to expose at its northbound some kind of control interface or protocol enabling the management of RTC media capabilities in a compatible way with the semantic requirements of the RTC Media API. The details about how to create such RTC media server and control protocol are out of the scope of this paper. The second, is to implement an RTC Media IDL compiler suitable for translating the RTC Media IDL into the corresponding programming-language-dependent SDKs. Remark that this compiler is not protocol

agnostic, in the sense that it needs to translate the RTC Media API invocations into the appropriate messages of the RTC media server control protocol. In other words, each specific media server control protocol needs to have a custom IDL compiler.

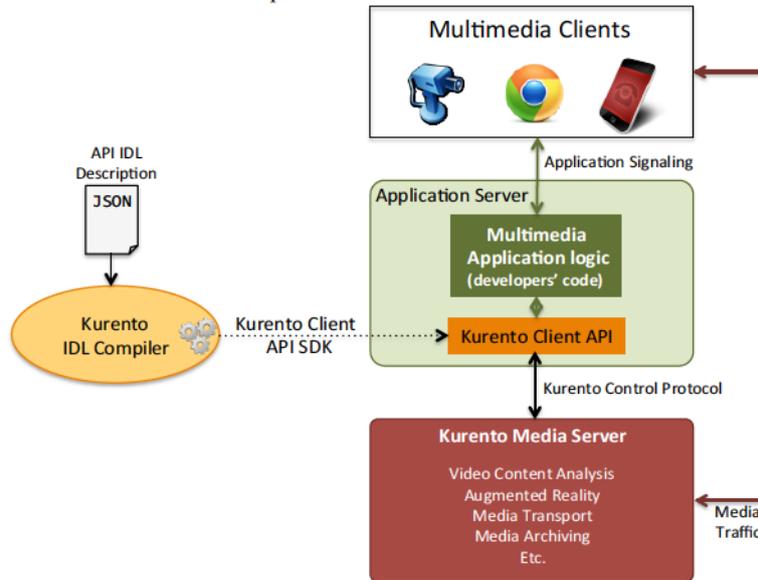


Fig. 5. Architecture of a Kurento application. As it can be observed, the Kurento IDL compiler generates the Java and JavaScript Kurento Client SDKs from the Kurento API IDL following the RTC Media API specifications described in this paper. Using these SDKs, developers can create their applications following the traditional three tiered WWW architecture just using the Kurento Client API as any other of their service APIs. The Kurento Client API implementation provides semantics to the API invocations by issuing the appropriate Kurento Control Protocol messages

In the context of the Kurento open source software project (<http://www.kurento.org>), we have created an example implementation of the RTC Media API. This implementation follows the architectural scheme depicted in Fig. 5. As it can be observed, our implementation provides the two above-mentioned ingredients. The Kurento Media Server plays the role of the RTC media server. Observe that Kurento Media Server exposes its capabilities through a JSON-RPC over WebSocket control protocol called the Kurento Control Protocol. This protocol has been designed to be compatible with the RTC Media API semantics. In addition, we have created an RTC Media IDL compiler capable of translating the IDL specifications into the appropriate API implementations both in Java and JavaScript. The resulting programming-language-dependent SDKs are called Kurento Client APIs in the Kurento jargon. Remark that the Kurento Client API is just a specific implementation of the RTC Media API suitable to interoperate with the Kurento Media Server.

Kurento is a complex technological stack and the interested reader can check the community documentation (<http://www.kurento.org/documentation>) and source code repositories (<https://github.com/kurento>) to have full information about the project. For the objectives of this paper, the interesting aspects are twofold. First, the Kurento Client API provides a full implementation of the RTC Media API as described in this paper. This implementation is specified through the RTC Media IDL and stored in files with the KMD.JSON extension (KMD for Kurento Module Description)². Second, the Kurento project provides a bunch of extensions to the RTC Media API in the form of custom filters, hubs and endpoints³. These extensions have been created following the RTC Media Module mechanism described above. Just for illustration, some of these extensions are presented here:

² Latest versions of the IDL files specifying the Kurento RTC Media API in the *kmd.json* format can be found at the Kurento GitHub repository at the following locations:

- <https://github.com/Kurento/kms-core/blob/master/src/server/interface>
- <https://github.com/Kurento/kms-elements/tree/master/src/server/interface>

³ Latest versions of the IDL files specifying the Kurento RTC Media API extensions in the *kmd.json* format can be found at the Kurento GitHub repository in different locations, including the following:

- <https://github.com/Kurento/kms-filters/tree/master/src/server/interface>
- <https://github.com/Kurento/kms-crowddetector/tree/master/src/server/interface>
- <https://github.com/Kurento/kms-platedetector/tree/master/src/server/interface>

- The `ZBarFilter` filter detects QR and bar codes in a video streams. When a code is found, the filter publishes a `CodeFoundEvent`. Application developers can add a listener to this event to execute some logic.
- The `ImageOverlayFilter` filter inserts still images in the video stream. The filter makes possible to select the position, scaling and rotation coordinates of the image.
- The `FaceOverlayFilter` filter detects faces in a video stream and overlays custom images onto the face coordinates. The filter makes possible to select specific scaling and offsets for the image position.
- `CrowdDetectorFilter` filter implements a computer vision algorithm suitable for detecting crowds of people into video streams. The level of crowdedness is published through a custom event that contains information about the direction and speed of movement of the crowd.
- `PlateDetectorFilter` filter detects European car plates and publishes the detected plate number as a custom event.
- `AugmentedRealityFilter` filter wraps the Alvar library [3] to provide marker and markerless Augmented Reality capabilities.
- `AlphaBlending` hub is a special type of hub that makes possible to mix different video streams using alpha transparency. This hub is useful for producing chroma blended videos in real time.

Thanks to all these capabilities, the Kurento software stack has been used for creating hundreds of applications combining different types features which include WebRTC and RTP transports, media recording, Video Content Analysis, or Augmented Reality⁴. All in all, Kurento provides a full working test-bed where the RTC Media API constructs described in this paper can be used, evaluated and improved.

3.1.9 Matching the RTC Media API requirements

To conclude with the RTC Media API presentation, we would like to come back to the list of requirements exposed in Section 2.3 to validate that they are fulfilled:

- **Seamless API extensibility through custom modules:** As it can be seen in Section 3.1.7, the RTC Media API can be extended in a seamless way by using the RTC Media Module mechanism, which provides full flexibility and no restrictions other than extending from the base RTC Media API classes.
- **Adaptation to WWW technologies and methodologies:** As shown in Section 3.1.8, RTC Media API implementations fully comply with the traditional WWW three tiered development model and enable developers to create applications leveraging novel WWW RTC media technologies such as WebRTC in a seamless and direct way.
- **Full abstraction of media details (i.e. codecs and protocols):** As it can be appreciated in the discussions and code examples in Sections 3.1.1, 3.1.3 and 3.1.7, the RTC Media API `connect` primitive exposed by all `MediaElements` makes possible to fully abstract codecs, protocols and formats. In all our examples there are no explicit references to codecs or formats, but many of the examples require specific transcodings to work. This is due to the fact the semantics of the `connect` primitive mandates the underlying media server capabilities to perform all the appropriate adaptations in a fully transparent way.
- **Programming language agnostic:** Discussions in Section 3.1.2 demonstrate the full agnosticism of the RTC Media API IDL in relation to programming languages. The only requirement for supporting a given programming language is to specify how the IDL is transformed into it and to implement the appropriate compiler following that specification. In Section 3.1.3 and 3.1.8 we provide such specifications and describe their implementations in Java and JavaScript in the context of the Kurento open source software project.
- **RTC media topology agnostic:** Following the discussions on Sections 3.1.1 and 3.1.6 one can appreciate that the RTC Media API makes possible to interconnect media elements following arbitrary and dynamic topologies thanks to the `connect` primitive. This means that developers do not need to be aware of the low level details of MMM, MSM or SFU technologies: they just need to interconnect their endpoints, filters and hubs accordingly to their needs. The RTC Media API semantics shall translate these interconnections into the appropriate low level mechanisms using MMMs, MSMs or SFUs in a fully transparent way.

⁴ Demos showing different Kurento applications creating using the Kurento RTC Media API implementations are accessible at the Kurento YouTube channel: <https://www.youtube.com/channel/UCFtGhWYqahVlzMgGntEmKug>

- **Advanced media QoS information gathering:** As it can be observed in the discussions in Section 3.1.2, the RTC Media API IDL does not restrict in any way the information a media object may expose through its properties and methods. We have leveraged such flexibility for creating QoS metrics gathering mechanisms in all endpoints based on the RTP protocol. In particular, the `WebRtcEndpoint` exposes primitives fully compliant with the standard WebRTC “inboundrtp” and “outboundrtp” stats [77].
- **Compatibility with advanced media processing capabilities:** As in can be observed in the discussions in Section 3.1.8, the Kurento software project has created a bunch of modules providing advanced capabilities such as Video Content Analysis, Augmented Reality, Computer Vision, etc. This demonstrates the ability of the RTC Media API `Filter` concept to hold all kind of extension for advanced media processing.
- **Context awareness:** The notion of context emerges in quite a seamless way from the discussions on Section 3.1.2. As it can be observed, the RTM Media API event mechanism makes possible for media capabilities to publish events to applications. These events may contain semantic information about the media content itself as shown, for example, in the `CrowdDetectorFilter` mentioned in Section 3.1.8. Hence, creating multimedia context-aware applications is straightforward: the application logic just needs to subscribe to the relevant events and publish them into a context database based on NGSI or any other equivalent standard.
- **Adapted to multisensory multimedia:** The RTC Media API can manage seamlessly arbitrary sensor data beyond audio and video. This can be achieved through the combination of two features. The first is the support for `DataChannel` that, as specified in Section 3.1.6, makes possible for any media pipeline to exchange multisensory multimedia with the external world using the WebRTC protocol stack. The second is the fact that, as described in Section 3.1.1, all streams exchanges among `MediaElements` may have a DATA track. In particular, any information received using `DataChannels` into a `WebRtcEndpoint` is published to the rest of the pipeline through the endpoint’s source DATA track. In the same way, any information received through the DATA track at a `WebRtcEndpoint`’s sink is sent to the network using `DataChannels`. As the `MediaElement` interface enables all the information received through the DATA to be used by the element internal logic, this mechanism makes possible, for example, to create Augmented Reality filters that leverage sensor information for customizing the augmentation logic.
- **Adaptation to cloud media servers:** As it can be observed in the discussion of Section 3.1.2, the RTC Media API does not specify how media pipelines are placed into media server instances. The API implementer has full freedom for selecting how newly created media pipelines are scheduled. This flexibility can be leveraged by API implementers to adapt their code to all kinds of cloud architectures. For example, as shown in the code snippet in Table 1, at the Kurento Client RTC Media API Java implementation, we decided that the RTC Media API is represented by a specific class (i.e. `KurentoClient`) that is built though a static `create` factory method. This method may accept as a parameter a single IP, in which case all pipelines are instantiated into the media server listening at that IP; a list of IPs, which causes media pipelines to be round-robin distributed on the corresponding media servers; or a media server scheduling interface, which can provide arbitrary logic for scheduling media pipeline creation into media servers. It may also accept no parameters and let the developer specify the behavior in a configuration file. All this flexibility makes possible for our RTC Media API to work seamlessly in cloud clusters of Kurento Media Server instances. Just as an example, this scheme is currently used in the NUBOMEDIA [49] and FIWARE [28] clouds. The complex details on how this happens are out of the scope of this paper. The point is to remark that the RTC Media API does not constraint the API implementer in any way when adapting to complex cloud scheduling and placement logic.

4 API evaluation

In the sections above we have presented the RTC Media API and we have introduced a specific implementation of it. The rest of this paper is devoted to describing a study we performed for evaluating the RTC Media API usability in the context of the Kurento open source software community.

4.1 Study design

For our study we used the CD framework [13]. Following common research practices [12], the evaluation of developers’ perception of our API is performed through a questionnaire. However, for the sake of clarity and simplicity, we included some adaptations. The CD framework define dimensions as a vocabulary that can be used by designers when investigating the cognitive implications of their design

decisions, so that designers might be able to express any properties of their information artifacts as a composition of these basic dimensions. As an analogy, this is somehow similar to the way vector spaces work: any vector in the space can be expressed as a composition of the base vectors. From this perspective, the base CD dimensions are designed for independence (i.e. they do not overlap) and not for clarity and simplicity. Due to this, the inclusion of the base CD dimensions in our questionnaire would make it excessively long, impractical and complex. This might decrease the aim of our target population to answer it and the overall usefulness of this research.

In this situation, we prefer to concentrate on an extended set of Clarke's dimensions [54], which are simpler and more intuitive. In particular, our analysis concentrates on 5 specific usability aspects: understandability, abstraction, expressiveness, reusability and learnability. All these high-level dimensions can be expressed in terms of the base CD dimensions, but their operationalization for our objectives is more practical because they shall allow to optimize the questionnaire for API users more than to API designers. Let's explain in details the meaning and value of each of our high-level dimensions.

Understandability deals with evaluating the effort required for understanding how to use the API for achieving a desired functionality. This dimension encompasses aspects such as whether the API names are descriptive and the relation among API types and constructs are clear and unambiguous. This relates to the base CD dimension called *closeness of mapping*. It also includes the ability of the API to avoid developers to manage hidden information not explicitly represented in the API, which is called *hidden information* in terms of the base CD dimensions. In addition, the base CD dimension called *hard mental operations* also affects understandability. In brief, this dimension addresses how simple is to access API features through object creation, primitive invocations or other means.

Abstraction, which is itself a base CD dimension, relates to the ability of the API to guarantee that programmers can use the API proficiency without requiring specific knowledge or assumptions in relation to its implementation details. Abstractions should match the conventions and practices of programmers, without being elegantly abstract at the expense of understandability or other practical concerns. Abstraction is typically correlated with the degree of comfort developers feel when using the API. Summarizing with a slogan, this research question asks whether the API "makes simple things simple, and complex things possible".

Expressiveness can be seen as the ability of inferring readily the purpose of an entity. This is related to the base CD dimension called *role-expressiveness*. Expressiveness is also related to how easy is for the programmer to build her code without needing to assume any specific cognitive model about API use. Intuitively, code written using expressive APIs tend to be simpler to read and transforming requirements into code is typically more efficient in expressive APIs. In terms of base CD dimensions these properties are related to *visibility* and *consistency*. We also consider the CD base dimension called *error-proneness* to be part of the expressiveness properties of our API.

Reusability determines whether the client code is maintainable and extensible. In particular, this dimension address the typical concern on how hard is to modify pre-existing code and adapt it to slightly, extended or more general requirements. The main related base CD dimensions is *viscosity*, understood as resistance to change, but it also involves other base dimensions such as *diffuseness* (i.e. the verbosity of the notation).

Learnability address the ability of the API learning process to be incremental. Learnable APIs enable developers to understand APIs in a gradual way without requiring initial disproportionate efforts, which is related to the CD base dimension called *progressive evaluation*. Learnability also deals with whether performing a certain programming task using the API has a positive impact on performing other related but different tasks. This dimension might have some overlap with understandability, but emphasizes specifically the learning process rather than its practical outcomes.

4.1.1 Research questionnaire

With the aim of enabling developers, rather than designers, to evaluate our API, we designed the questionnaire in a way that end users can readily understand. For this, we adapted the commonly accepted survey model in the CD community [12] to our specific high-level dimensions and audience. All in all, the questionnaire consists of 28 assertions which characterize the 5 target dimensions. For every assertion, users provide their degree of agreement or disagreement in a scale from 1 (I fully disagree) to 5 (I fully agree). For assessing the internal consistency of the data, and following common practices in

psychological research, some of the assertions are generated in negative terms. For example, if a responded expressed agreement with the claim “I feel this API is simple” and disagreement with “I find it’s hard programming with this API” this would be an indication of internal consistency. We call these assertions and N-assertions. For the statistical analysis, answers of N-assertions are inverted (i.e. 1 is transformed into 5, 2 into 4, 4 into 2 and 5 into 1), so that consistency and coherence is maintained.

In addition to this, and with the objective of characterizing different aspects of participants, a number of questions were included for profiling demographic data and for evaluating their degree of experience. These questions are shown on Table 17. As it can be observed, most of the questions are self-explanatory with these exceptions:

- For the question “Type of application being developed”, we provided participants the possibility of selecting multiple items among the following options (the corresponding encoding token is provided between < > signs):
 - o I’m creating video applications with recording capability (<Recording>)
 - o I’m creating my own filters and extending Kurento APIs (<Filter>)
 - o I’m creating video surveillance applications (<Video surveillance>)
 - o I’m creating videoconferencing applications (<Videoconferencing>)
 - o I’m creating broadcasting applications for distributing media among large groups of receivers (<Broadcasting>)
 - o I’m using Kurento for integrating with other types of technologies beyond WebRTC (<Integration>)
 - o I’m creating other types of applications (<Other>)
- All questions dealing with “self-assessment of expertise” were expressed in the poll as assertions in the form “I’m an expert in ...”, where answers are in the above mentioned 1 (I fully disagree) to 5 (I fully agree) format.
- The question dealing with “Learning stage on Kurento technologies” provided the ability of selecting one option among the following items:
 - o I tried to install Kurento unsuccessfully
 - o I installed Kurento and executed some of the provided demos
 - o I developed a simple application
 - o I developed a complex application
 - o I developed a complex application which is in production

Table 16. Research questionnaire used for evaluating developer’s perception of API usability on the 5 target dimensions. Every assertion in the questionnaire is identified with a unique ID for further reference (e.g. U.1 refers to the first question of the Understandability dimension). Assertions generated in negative terms (i.e. N-assertions) start with an (N) mark. These assertions are useful for evaluating the consistency of the research. Participants are asked to provide their degree of agreement with every assertion in a scale from 1 (I fully disagree) to 5 (I fully agree). For the statistical analysis, N-assertions and inverted so that the coherence of the questionnaire is maintained

Dimension	ID	Assertion
Understandability	U.1	Kurento APIs are, in general, easy to understand
	U.2	In Kurento APIs object names are descriptive and unambiguous
	U.3	(N) I need to keep track of hidden information not represented by the APIs to create my applications
	U.4	(N) Kurento APIs are obscure and it takes a huge effort to use them, even for creating simple applications
	U.5	Kurento API objects, types, and primitives represent appropriately the underlying media-related concepts
	U.6	I understand the difference between a Media Pipeline and a Media Element
Abstraction	A.1	Kurento APIs make simple to create applications without needing to worry about the low level media details
	A.2	(N) I needed to adapt the API (e.g. inheriting, overriding, etc.) for having it meet my needs
	A.3	(N) It’s necessary to understand how codecs and protocols work for being able to use Kurento APIs
	A.4	I like writing applications with Kurento APIs. I’m familiar with their programming model
	A.5	I feel appealing and attractive the general approach of Kurento APIs
	A.6	Creating simple applications with Kurento APIs is simple. Creating complex applications is possible
Expressiveness	E.1	Developing with Kurento fully matches the expectations I had
	E.2	I can translate my media application requirements into code in an easy way
	E.3	Reading an application code, I can understand what the application is doing in a simple way
	E.4	After creating an application, I can explain seamlessly to other people what I have done in terms of media elements and their interconnections
	E.5	(N) There are missing features in Kurento APIs that make not possible to implement interesting applications
	E.6	(N) Programming with Kurento APIs is error prone. You need to take into consideration a lot of details for having an application working
Reusability	R.1	(N) Creating applications requires too long and verbose code specifying too many things
	R.2	(N) Adding a recording capability to an non-recording application requires modifying a lot of code

	R.3	My code using Kurento APIs can be maintained and evolved easily
	R.4	I can re-use Kurento related code in a simple way
	R.5	(N) When using Kurento APIs, there are many different ways of doing the same thing and I need to take too many decisions in the process
	R.6	Adapting my Kurento-based application to new media requirements is quite simple
Learnability	L.1	I learned how to use Kurento in an incremental way, starting with simple concepts and progressing towards complex applications
	L.2	(N) Programming with Kurento requires learning a lot of classes and dependencies, even for applications.
	L.3	When I create a complex application, I can start by simple example and evolve it later in a seamless way
	L.4	(N) I needed to read all the documentations and tutorials to be able to create my first application
	L.5	Reading simple tutorials made me possible to understand better the complex ones and to create later applications complying with my requirements

4.1.2 Participants and protocol

Most API usability studies are performed by recruiting students or researchers who are trained on the API through lectures or exercises and who are later interviewed for performing the evaluation [54]. These types of protocols are sensitive to many different types of bias that may affect the study reliability. In particular, their main weakness is that participants are typically not professional developers and are not faced with real world programming tasks. Hence, their perception of the API limitations and usability problems can be severely biased by their own background and by the nature and contents of the training contents and of the proposed exercises. In addition, those contents and exercises are typically created by the API designers, which increases significantly the risk of introducing the designers' cognitive models and hiding API limitations that might not be known even by designers themselves. Many API evaluation research works are aware of these limitations, but solving them is not trivial given the difficulty of reaching a statistical significant population of professional developers being independent of designers, having the time of learning API concepts and working on solving real-world tasks with them.

Table 17. This table shows the additional questions made to participants in order to characterize them. The first column shows the type of data to be gathered through the question. The second column shows the question itself (summarized for the sake of readability), the third column shows the value type accepted by the web form. The mark [] indicates users are given the choice of choosing one item in a list. The mark []* indicates that multiple items may be selected. In this table, list items are tokenized for simplicity

Question type	Question	Value type
Demographic data	Nationality	[Country select]
	Age	Integer
	Gender	[Male Female]
	Main computer language used	[Java JavaScript Other]
	Type of application being developed	[Recording Filter Video surveillance Videoconferencing Broadcasting Integration Other]*
Development experience	Years of experience as developer	Integer
	Hours learning or programming with Kurento	Integer
	Self-assessment of expertise as WebRTC developer	[1 to 5]
	Self-assessment of expertise in media technologies	[1 to 5]
	Self-assessment of expertise with Kurento technologies	[1 to 5]
	Learning stage on Kurento technologies	[Tried Installed SimpleApp ComplexApp ProductionApp]

In order to avoid these problems, we leverage the fact that the RTC Media API has been implemented as part of the Kurento project. More specifically, the Kurento Client API is an almost complete implementation of it. This is a significant advantage because Kurento has been released as an Open Source Software and a community of developers has emerged around it. The size of the community is unknown but its main communication channel, the Kurento Public mailing list, has, at the time of this writing, 432 subscribers most of which are professional developers at different stages of the API learning process.

In this context, the survey protocol is simple. The questionnaire is designed for Kurento community members (i.e. in all assertions the API is referred to as the "Kurento API"). This questionnaire is published as a web form and participants are invited to participate through an aseptic e-mail invitation sent to the Kurento Public mailing list. This mail is written trying to avoid any kind of bias on participants so that it just presents the survey objectives and exposes a privacy policy guaranteeing that no personal data is to be disclosed or used for other objectives than the ones of the survey. Participants are incentivized to participate only once by requiring logging into the web form system through a valid e-mail. The form makes mandatory to fill answers on all assertions for making submission possible (i.e.

partially answered questionnaires are not considered). The web form system stores in a persistent database each participant's answers and makes possible to edit them during the duration of the survey period, which is limited to 2 weeks' time.

4.2 Results and discussion

4.2.1 Analysis of participants

The survey was activated following the protocol described above. After one week from the initial e-mail invitation a total of 17 participants had answered. Several reminders were sent to the Kurento Public mailing and the announcement was also published through different social channels, such as the Kurento Twitter account. In two weeks, 42 answers were received, which represents 9.7% of the number of Kurento Public mailing list subscribers. This is aligned with typical answer rates in surveys.

From a demographic perspective age of participants were distributed between 20 and 50 years old, being the more numerous group the one in the thirties, which account for 50% of the total participants, as can be seen in Fig. 6. It is worth noting that 100% of participants were male, and their main programming language, the language in which API usage was involved, was JavaScript, totaling a 62% of respondents.

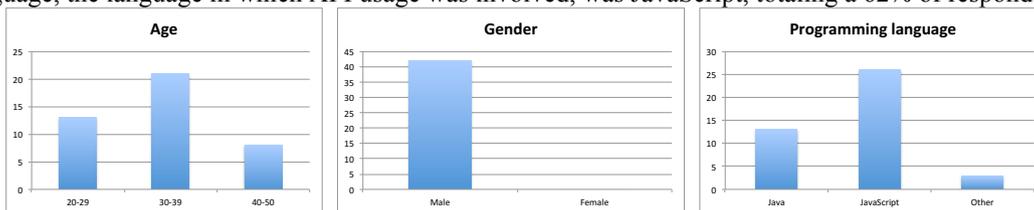


Fig. 6. Figures showing the total number of participants for different demographic data including Age (left), Gender (middle) and Main Programming Language (right)

In relation to nationality, and as can be seen in Fig. 7, the poll was answered by developers from 20 nationalities on 4 different continents, being USA the country with more participants.

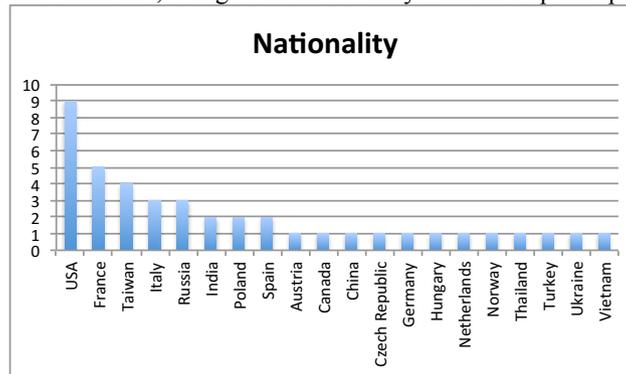


Fig. 7. Total number of participants per nationality

In Fig. 8 we show the types of applications being developed. WebRTC video broadcasting applications (for distributing a media stream among a group of developers) and videoconferencing services are the most popular with 27% and 26% respectively. Applications involving recording and WebRTC integration are also quite popular (19% and 15%). Media processing services, such as video surveillance applications and services requiring custom filters are less popular accounting only for 7% and 5% respectively.

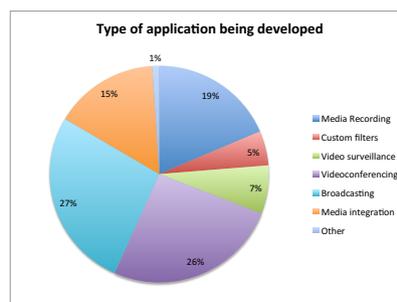


Fig. 8. Types of applications being developed with the Kurento API. The classification is based on the types of consumed features. In the poll, developers were able to select several classes of features for their applications

When coming to expertise evaluation, as shown in Table 18, participants accumulate, on average, 9.7 years of development experience, but with a relevant diversity (i.e. 1 year as minimum and 30 as maximum). Participants also declare to have invested, on average, more than 61 hours for learning and programming with Kurento technologies, which provides a reasonable guarantee on the their ability to evaluate the API. Regarding the self-assessment of competences, a relevant expertise on WebRTC technologies is clearly declared (mean 2.6 and median of 3). However, participants seem to have more uncertainty in relation to their knowledge of general video technologies (mean 2.6 and median of 2) and Kurento technologies (mean 2.5 and median of 2). On the latter, no participants declare to feel as a full featured expert (maximum is 4).

Table 18. This table shows the summary of answers in relation to participant’s expertise as developers and in the different involved technological areas

Expertise question	Mean	Median	Min	Max	STDDEV	Type of data
Years of experience as developer	9.7	5	1	30	8	Integer
Hours learning or programming with Kurento	61.6	20	2	500	127	Integer
Self-assessment on WebRTC expertise	2.7	3	1	5	0.9	1 to 5
Self- assessment on video technologies expertise	2.6	2	1	5	1.2	1 to 5
Self- assessment on Kurento expertise	2.5	2	1	4	1.1	1 to 5

In relation to the question regarding the stage of learning on Kurento technologies an interesting surprise emerges given than 38% of participants declare to have already an application in production, while 19% and 33% claim to have developed one complex and one simple application respectively. Only 3% of participants has not been able to install and test Kurento and its APIs. These results are shown in Fig. 9.

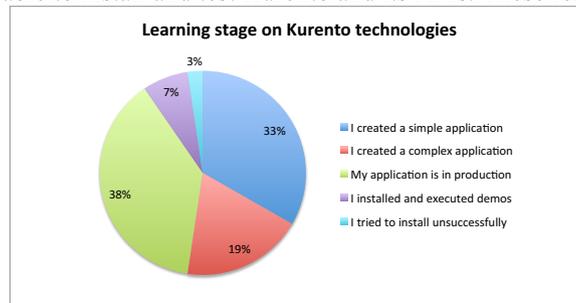


Fig. 9. Pie chart showing the level of expertise in the Kurento technologies of the participants

4.2.2 Analysis of dimensions

The results of the poll are summarized in Table 19, where the main statistics for each of the assertions and their corresponding dimensions are depicted. As specified above, all N-assertions were inverted previous to the statistical analysis. Hence, magnitudes represent perceptions on API usability in positive terms (i.e. the higher the magnitude the better the developer’s impression on the APIs). As it can be seen in Fig. 10, on average, participants feel API usability properties are adequate, being reusability the dimension with highest rank and expressiveness the one with lower score. A detailed analysis for each dimensions is presented in the following paragraphs.

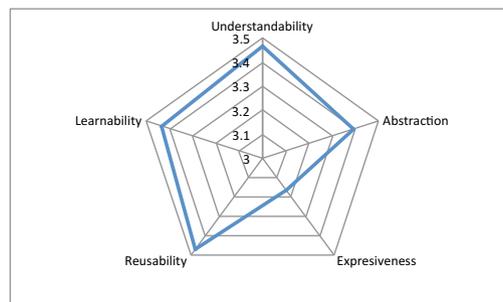


Fig. 10. Radar chart showing average rankings on the 5 target dimensions of our questionnaire. Scale is set between 3 and 3.5 for evidencing the differences among dimensions

Understandability.

As a general perception, participants feel the API understandability is fine, with an average of 3.46 over all answers on this dimension. Through U.1 (3.33) we find a general declaration that the API is easy to understand. In particular, and as shown in assertion U.2 answers, participants feel outstanding (4.26) how

object and primitive names are descriptive. On the other hand, and through the U.3 N-assertion (2.69), we find that developers detect the presence of hidden dependencies that should be avoided.

Abstraction.

When coming to abstraction, we also find out a general positive evaluation (the overall average is 3.39). Answers on all questions are quite uniform, being A.5 the top ranked one (3.69), which shows that developers feel appealing the API approach; and A.2 the bottom ranked one (3.24) indicating that some developers feel the need of adapting the API to their needs. It is remarkable that A.2 has the largest standard deviation (1.28) of A assertions, which evidences some degree of controversy. This is confirmed when looking closely to answers: in A.2 answers, there are 12% of ones and 19% of fives, while considering all Abstraction answers the ratios for ones and fives are 4.3% 12.6 respectively.

Expressiveness.

The Expressiveness analysis also reflects positive evaluation but shows improvement areas. This is the less successful dimension with an overall average ranking of 3.17. Expressiveness limitations seem to emerge on assertions U.5 and U.6 (2.76 and 2.90 respectively). In particular, U.5 reveals that developers miss features that are relevant for their applications. U.6, in turn, manifests what in CD terminology is called *error proneness* or, in other words, that the API gives not enough protection against failures. Both issues might be somehow explained by the inherent complexity of RTC technologies. In any case, this is clearly a topic that deserves further investigation. On the other hand, as demonstrated through U.3 and U.4, our API has suitable visibility for code reading (3.43) and is consistent when dealing with explaining code logic in terms of the API constructs.

Reusability.

Reusability is the dimension with the highest ranking (3.47 in average). Our API seems to have nice properties in relation to what, in CD jargon, is called *viscosity* (i.e. the resistance for code consuming the API to be modified and evolve). This is illustrated through results in assertions U.2, U.3 and U.6, which average to 3.67, 3.31 and 3.14 respectively. Also the API demonstrates nice properties in relation to what's called *diffuseness* in the CD literature as shown by the 3.48 exhibited by U.1.

Learnability.

As with the rest of our dimensions, learnability is evaluated positively by participants (3.43 in average). The improvement area in this topic emerges from L.2 (3.07) and L.4 (2.86), which evidence the impressions of developers of needing to learn about a lot of API constructs and to read a relevant amount of documentation before being able to start using the API for useful things. On the other hand, as U.1 (3.88) evidences, the learning process seems to be compatible with an incremental approach where complexity is introduced in progressive steps.

Table 19. Results of the research showing, for each assertion of the poll, the main statistics of the provided answers. Notice that N-assertion have their results inverted for maintaining coherence. For each dimension, the statistics are computed on the average values along all assertions on that dimension for each user

Dimension	ID	Mean	Median	Min	Max	STDDEV
Understandability		3.46	3.58	2	4.67	0.58
	U.1	3.33	3	1	5	0.98
	U.2	4.26	4	2	5	0.86
	U.3	2.69	3	1	5	1.12
	U.4	3.26	3	1	5	1.17
	U.5	3.52	3.5	2	5	0.71
	U.6	3.71	4	1	5	1.20
Abstraction		3.39	3.42	1.83	4.23	0.57
	A.1	3.40	3.50	1	5	1.04
	A.2	3.24	3	1	5	1.28
	A.3	3.29	3	1	5	1.11
	A.4	3.26	3	1	5	0.86
	A.5	3.69	4	1	5	0.84
	A.6	3.48	4	2	5	0.92
Expressiveness		3.17	3.17	2	4.33	0.57
	E.1	3.29	3	2	5	0.71
	E.2	3.17	3	1	5	0.82
	E.3	3.43	4	1	5	0.97
	E.4	3.45	4	1	5	1.02
	E.5	2.76	3	1	5	1.19
	E.6	2.9	3	1	4	1.05
Reusability		3.47	3.5	1.67	4.67	0.52
	R.1	3.48	4	1	5	0.99
	R.2	3.67	4	1	5	1.16
	R.3	3.31	3	1	5	0.98
	R.4	3.6	4	1	5	0.7
	R.5	3.62	4	1	5	0.7
	R.6	3.14	3	1	4	0.75
Learnability		3.43	3.4	1.4	5	0.65
	L.1	3.88	4	2	5	0.77
	L.2	3.07	3	1	5	1.35
	L.3	3.67	4	2	5	0.95
	L.4	2.86	3	1	5	1.26
	L.5	3.69	4	1	5	0.95

To complete our discussion, we perform an additional analysis to assess how participant experience might be influencing API perception. In this sense, we compute the average of all the scores provided by participants along all assertions and calculate its correlation coefficient with the available experience-related variables. Table 20 show the numerical results of this analysis, which evidence that higher development experience tends to be associated with better usability perception. This is particularly true for the self-assessment expertise on Kurento technologies, which is not surprising: developers having better perception on their expertise on Kurento feel the Kurento API to be more usable for their objectives. However, as Fig. 11 shows, the negative effect is concentrated on users with very low expertise on Kurento technologies. In other words: as soon as a developer feels to have some initial knowledge on the API, his perception of the usability increases to be the same as the one of experts, which is a good symptom. Interestingly enough, this effect seems not to be correlated with the number of hours declared in learning or programming with Kurento. This means that the time invested by developers in learning or programming with the API does not have a strong influence on their perception of API usability. This may be related with the fact that the self-assessment of expertise on WebRTC technologies and on video technologies are much correlated with the self-assessment of expertise on Kurento (the correlation coefficient is 0.55 and 0.34 respectively). This evidences that having a previous understanding on WebRTC and video technologies enables our API developers to improve their API usability perception in a faster and seamless way.

Table 20. This table shows the correlation between the different parameter data captured through the questionnaire against the scores of API usability perception averaged across all assertions

Parameter to be correlated with overall average score along all answers	Correlation
Years of experience as developer	0.16
Hours learning or programming with Kurento	-0.07
Self-assessment of expertise as WebRTC developer	0.43
Self-assessment of expertise in media technologies	0.22
Self-assessment of expertise with Kurento technologies	-0.04

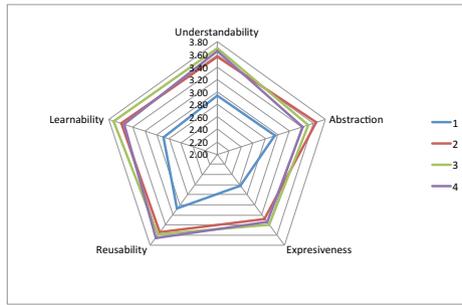


Fig. 11. Radar chart showing the average API usability for each of the analyzed dimensions particularized for the 4 levels of expertise declared for the question “Self-assessment of Kurento expertise” (notice that no participants declared a value of 5). As it can be observed, users declaring very low level of expertise (i.e. 1) tend to perceive the API as less usable in all dimensions. However, users declaring 2 or more perceive the API in quite a similar way

To conclude, we also evaluated the correlation between the API usability perception and the main programming language being used. Results are summarized in Fig. 12. As it can be observed, JavaScript developers have a better perception of API usability than Java developers, which might evidence that our API ideas might be slightly better adapted to the idiosyncrasy and methodology of JavaScript. Interestingly, developers in the “Other” category have much degraded perception of the API usability, significantly in what refers to understandability. This may cause by the fact that the only official implementations of the Kurento Client API are the Java and JavaScript ones, which means that the “Other” developers are either working with non-official API implementations or directly consuming the JSON-RPC over WebSocket protocol exposed by Kurento Media Server, which enables the same type of functionalities of the wrapper Java and JavaScript APIs, but which is clearly more complex and cumbersome to be used.

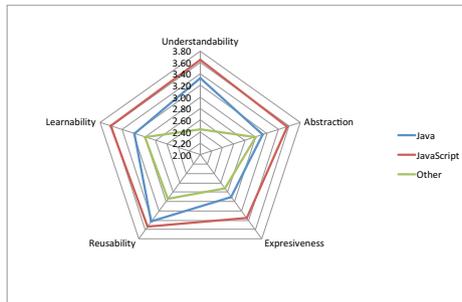


Fig. 12. Radar graph showing the average API usability perception as a function of the main programming language being used: Java, JavaScript or Other. The axis are re-scaled for clarity. Remark that the Kurento Client API is only available in Java and JavaScript. Hence, users in the category “Others” are directly consuming the JSON-RPC over WebSocket protocol exposed by Kurento Media Server, which is significantly more complex

4.3 Validity of the analyses

Following commonly accepted techniques for evaluating assessment data [22], we discuss the main threads to the validity of our research as well as the measures we deployed to minimize their impact.

4.3.1 Construct validity

Construct validity is the degree to which a test measures what it claims, which in our case, is the degree of usability of our RTC Media API. As introduced in sections above, evaluating API usability is a very complex multifaceted problem where it is difficult to get objective measures. To minimize threads, we performed a careful design of our research protocol which included the following protection mechanisms:

- We used a well-established methodology based on the CD framework, the most widely accepted technique for this objective which has been already used successfully on a number of usability studies worldwide.
- We performed a careful design of the questionnaire basing on high-level usability dimensions more adapted to participants needs than to API designer’s needs. Each of the high-level dimensions was measured through groups of 5 to 6 assertions facing the problem from different perspectives which minimizes effects of assertion misinterpretation.
- The questionnaire contained complementary questions digging into the different components of each of the high-level dimensions and combining positively and negatively formulated assessments. This should enhance the consistency guarantees of the answers.
- The protocol avoided to introduce any kind of bias by enabling participants to answer assertions basing only on their own knowledge about the API artifacts (i.e. documentation, code, etc.) and not

on previous information provided by designers (e.g. training courses) or on specific artificial exercises which could be associated to specific cognitive models about the API.

4.3.2 Internal validity

Internal validity is a property associated to the extent to which a study minimizes systematic errors and avoid introducing bias into measurements. For enhancing our internal validity, in our research we tried to avoid any kind of selection bias by enabling Kurento Open Source Community members to answer freely to the poll. This strategy was clearly successful given the wide spectrum of participants we had, comprising developers of different ages, expertise degrees, nationalities and cultures. This enhances significantly our internal validity in relation to other previous similar studies [54] where API designers and participants have tight relationships (e.g. professors and students, workers of the same company, etc.) The risk of statistical effects in the data is also low given the fact that the poll was answered by 42 participants, which is a population sample significantly over the ones of other similar studies [54].

To formalize our internal validity analysis, we perform an additional test based on Cronbach's alpha [18, 19], which is the most commonly used estimate for assessing the reliability of psychometric tests in social sciences. As Cronbach's alpha is a measure of the internal consistency of data it needs all test items to measure the same construct. Due to this, its computation needs to be performed for each of our high-level dimensions separately. Cronbach's alpha evaluation for our data is shown in **Table 21**, while the thresholds for interpreting it are depicted in Table 22.

Table 21. Cronbach's alpha computed for all the high-level dimensions of our test

Dimension	Cronbach's alpha
Understandability	0.77
Abstraction	0.73
Expressiveness	0.76
Reusability	0.76
Learnability	0.71

Table 22. Commonly accepted rule of thumb for describing internal consistency in terms of Cronbach's alpha

Cronbach's alpha	Internal consistency
$\alpha \geq 0.9$	Excellent
$0.9 > \alpha \geq 0.8$	Good
$0.8 > \alpha \geq 0.7$	Acceptable
$0.7 > \alpha \geq 0.6$	Questionable
$0.6 > \alpha \geq 0.5$	Poor
$0.5 > \alpha$	Unacceptable

As it can be observed, the reliability of the obtained data is inside acceptable margins, which is reasonable for our type of questionnaire and open research search methodology where there is no control on the who's, how's and why's of participants.

4.3.3 External validity

External validity refers to the extent to which a study can be generalized to other situations or populations. In relation to this, the main threads to external validity come from the research protocol, which was designed basing on specificities of our API implementation. In particular, the fact that we leveraged the Kurento Open Source software community for obtaining the test population is quite a strong restriction for the generalizability of our findings given that most newly designed APIs might not be Open Source and, even if they are, they do not need to have an active international community of more than 400 developers. Otherwise, the rest of our methodology as well as the analysis we performed and the conclusions we gathered from it do not assume any kind of specific requirements. This suggests that the gist of our findings are also applicable to different contexts and populations.

5 Conclusions

In this paper we have presented the RTC Media API: a new type of application programming interface complying with a number of stringent requirements answering to latest trends and needs in the area of real-time multimedia. We have introduced a specification of the API showing how it can be formally defined in a programming-language agnostic way through an IDL. We have also demonstrated how the IDL can be compiled into different programming languages, including Java and JavaScript, while complying with latest trends on API usability and performance. We have presented a specific RTM Media API implementation, the Kurento Client API, which includes a bunch of advanced processing capabilities such as Video Content Analysis, Augmented Reality, Media Blending, etc. To conclude, we have evaluated the API usability through a research study basing on the CDs framework to demonstrate

the API suitability in terms of dimensions such as understandability, abstraction, expressiveness or learnability.

Along the paper, we have tried to stress the importance of listening to developers' needs and solving developers' problems. We defend software is eating RTC multimedia technologies. Hence, for pushing them to a next level we need to create novel APIs and SDKs suitable for their democratization among wider developer audiences. In current state-of-the-art, there are a huge amount of algorithms and technologies for transporting, analyzing and enriching media, but there are very few APIs and SDKs making possible for average WWW and smartphone developers to use them in a seamless and effortless way. Our RTC Media API brings a whole new concept by incorporating WWW development methodologies to the multimedia arena.

The RTC Media API in general, and the Kurento Client API implementation in particular, are still research artifacts, which are under maturation, and miss many relevant ingredients. In particular, the adaptation to latest trends in WebRTC technologies, including the incorporation of ORTC (<http://ortc.org/>) concepts, would benefit significantly the API flexibility and its ability to adapt to latest trends in WebRTC. The API would also benefit from having richer support for complex media streams suitable for including multiple audio and video tracks, so that 3D or MVC multimedia is supported. Improvements are also possible from the perspective of development tools beyond the API itself so that seamless mechanisms for debugging, diagnosing and optimizing applications would be more than welcome by developers.

Acknowledgments This work has been supported by the European Commission under projects FI-WARE FP7-2011-ICT-FI GA-285248, FI-CORE FP7-2014-ICT-FI GA- 632893 and NUBOMEDIA FP7-ICT-2013-1.6 GA-610576; by Spanish Ministerio de Educación under project Reactive Media (TIN2013-41819-R); and by the Regional Government of Madrid (CM) under project Cloud4BigData (S2013/ICE-2894) co-funded by FSE & FEDER.

References

1. Abowd GD, Dey AK, Brown PJ, Davies N, Smith M, Steggles P (1999) Towards a better understanding of context and context-awareness. In: *Handheld and ubiquitous computing*, Springer, pp 304-307
2. Afonso LM, Cerqueira RFdG, de Souza CS (2012) Evaluating application programming interfaces as communication artefacts. *System* 100:8-31
3. Ajanki A, Billingham M, Gamper H, Järvenpää T, Kandemir M, Kaski S, Koskela M, Kurimo M, Laaksonen J, Puolamäki K, et al (2011) An augmented reality interface to contextual information. *Virtual reality* 15(2-3):161-173
4. Alvestrand H (2015) Transports for WebRTC draft-ietf-rtcweb-transports-10. <https://tools.ietf.org/html/draft-ietf-rtcweb-transports-10>, accessed 20 December 2015
5. Andreasen F, Arango M, Huitema C, Kumar R, Pickett S, Elliott I, Foster B, Dugan A (2003) Media gateway control protocol (MGCP) version 1.0. Tech. rep., Internet Engineering Task Force, Request For Comments (RFC) 3435
6. Bainomugisha E, Carreton AL, Cutsem Tv, Mostinckx S, Meuter Wd (2013) A survey on reactive programming. *ACM Comput Surv* 45(4):52:1-52:34, DOI 10.1145/2501654.2501666, URL <http://doi.acm.org/10.1145/2501654.2501666>
7. Andreessen M (2011) Why software is eating the world. *Wall Street Journal* 20
8. Bajaj V (2004) JAIN MEGACO API Specification. Tech. rep., Java Community Process, Java Specification Request (JSR) 79
9. Bauer M, Kovacs E, Schülke A, Ito N, Criminisi C, Goix LW, Valla M (2010) The context API in the oma next generation service interface. In: *Intelligence in Next Generation Networks (ICIN)*, 2010 14th International Conference on, IEEE, pp 1-5
10. Becke M, Rathgeb EP, Werner S, Rungeler I, Tuxen M, Stewart R (2013) Data channel considerations for rtweb. *Communications Magazine*, IEEE 51(4):34-41
11. Bertoa MF, Troya JM, Vallecillo A (2006) Measuring the usability of software components. *Journal of Systems and Software* 79(3):427-439
12. Blackwell AF, Green TR (2000) A cognitive dimensions questionnaire optimised for users. In: *Proceedings of the Twelfth Annual Meeting of the Psychology of Programming Interest Group*, pp 137-152

13. Blackwell AF, Britton C, Cox A, Green TR, Gurr C, Kadoda G, Kutar M, Loomes M, Nehaniv CL, Petre M, et al (2001) Cognitive dimensions of notations: Design tools for cognitive technology. In: *Cognitive Technology: Instruments of Mind*, Springer, pp 325-341
14. Bloch J (2006) How to design a good API and why it matters. In: *Companion to the 21st ACM SIGPLAN symposium on Object-oriented programming systems, languages, and applications*, ACM, pp 506-507
15. Bray T, Paoli J, Sperberg-McQueen CM, Maler E, Yergeau F (1998) Extensible markup language (XML). World Wide Web Consortium Recommendation REC-xml-19980210 <http://www.w3.org/TR/1998/REC-xml-19980210>
16. Catherine MR, Edwin EB (2013) A survey on recent trends in cloud computing and its application for multimedia. *International Journal of Advanced Research in Computer Engineering & Technology (IJARCET)* 2(1):304-309
17. Clarke S (2001) Evaluating a new programming language. In: *13th Workshop of the Psychology of Programming Interest Group*, pp 275-289
18. Cortina JM (1993) What is coefficient alpha? An examination of theory and applications. *Journal of applied psychology* 78(1):98
19. Cronbach LJ (1951) Coefficient alpha and the internal structure of tests. *Psychometrika* 16(3):297-334
20. Daughtry III JM, Carroll JM (2012) Perceived self-efficacy and APIs. *Programming Interest Group* p 42
21. Dey AK (2001) Understanding and using context. *Personal and ubiquitous computing* 5(1):4-7
22. Downing SM (2003) Validity: on the meaningful interpretation of assessment data. *Medical education* 37(9):830-837
23. Duala-Ekoko E, Robillard MP (2012) Asking and answering questions about unfamiliar APIs: An exploratory study. In: *Proceedings of the 34th International Conference on Software Engineering*, IEEE Press, pp 266-276
24. Ellis B, Stylos J, Myers B (2007) The factory pattern in API design: A usability evaluation. In: *Proceedings of the 29th international conference on Software Engineering*, IEEE Computer Society, pp 302-312
25. Ericson T, Brandt M (2009) Media Server Control API. Tech. rep., Java Community Process, Java Specification Request (JSR) 309
26. Farooq U, Zirkler D (2010) API peer reviews: a method for evaluating usability of application programming interfaces. In: *Proceedings of the 2010 ACM conference on Computer supported cooperative work*, ACM, pp 207-210
27. Ferry D, Lim S (2004) JAIN SLEE API Specification. Tech. rep., Java Community Process, Java Specification Request (JSR) 22
28. FIWARE Consortium (2015) Future Internet Core Platform. Available online: <http://www.fiware.org>, FP7-2011-ICT-FI (GA-285248)
29. Fraternali P (1999) Tools and approaches for developing data-intensive web applications: a survey. *ACM Computing Surveys (CSUR)* 31(3):227-263
30. Gamma E, Helm R, Johnson R, Vlissides J (1994) *Design Patterns. Elements of Reusable Object-Oriented Software*. Addison-Wesley.
31. Gouveia F, Wahle S, Blum N, Magedanz T (2009) Cloud computing and EPC/IMS integration: new value-added services on demand. In: *Proceedings of the 5th international ICST mobile multimedia communications conference*, ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering), p 51
32. Green TR (1989) Cognitive dimensions of notations. *People and computers V* pp 443-460
33. Green TR, Blandford AE, Church L, Roast CR, Clarke S (2006) Cognitive dimensions: Achievements, new directions, and open questions. *Journal of Visual Languages & Computing* 17(4):328-365
34. Gualdi G, Prati A, Cucchiara R (2008) Video streaming for mobile video surveillance. *Multimedia, IEEE Transactions on* 10(6):1142-1154
35. Handley M, Perkins C, Jacobson V (2006) SDP: session description protocol. Tech. rep., Internet Engineering Task Force, Request For Comments (RFC) 4566
36. Henning M (2007) API design matters. *Queue* 5(4):24-36
37. Hovemeyer D (2005) Simple and effective static analysis to find bugs. PhD thesis, College Park, MD, USA, AAI3184274
38. Internet Engineering Task Force (2015) Real-time communication in web-browsers (rtcweb). <https://datatracker.ietf.org/wg/rtcweb>, accessed 9 December 2015

39. Ivov E, Marocco E, Holmberg C (2014) A Session Initiation Protocol (SIP) usage for Trickle ICE draft-ietf-mmusic-trickle-ice-sip-03. <https://tools.ietf.org/html/draft-ietf-mmusic-trickle-ice-sip-03>, accessed 12 December 2015
40. Jennings C, Holmberg C, Alvestrand HT (2015) Negotiating media multiplexing using the session description protocol (SDP) draft-ietf-mmusic-sdp-bundle-negotiation-23. <https://tools.ietf.org/html/draft-ietf-mmusic-sdp-bundle-negotiation-23>, accessed 12 December 2015
41. Johnston AB, Burnett DC (2012) WebRTC: APIs and RTCWEB protocols of the HTML5 real-time web. Digital Codex LLC
42. Kambona K, Boix EG, De Meuter W (2013) An evaluation of reactive programming and promises for structuring collaborative web applications. In: Proceedings of the 7th Workshop on Dynamic Languages and Applications, ACM, New York, NY, USA, Dyla '13, pp 3:1-3:9, DOI 10.1145/2489798.2489802, URL <http://doi.acm.org/10.1145/2489798.2489802>
43. Koukoulidis V, Shah M (2006) The IP multimedia domain: service architecture for the delivery of voice, data, and next generation multimedia applications. *Multimedia Tools and Applications* 28(2):203-220
44. Kristensen A (2003) SIP Servlet API. Tech. rep., Java Community Process, Java Specification Request (JSR) 116
45. Laursen A, Olkin J, Porter M (1994) Oracle media server: Providing consumer based interactive access to multimedia data. *SIGMOD Rec* 23(2):470-477, DOI 10.1145/191843.191933
46. Lienhart R, Maydt J (2002) An extended set of HAAR-like features for rapid object detection. In: *Image Processing. 2002. Proceedings. 2002 International Conference on*, vol 1, pp I-900-I-903 vol.1, DOI 10.1109/ICIP.2002.1038171
47. Melanchuk T (2009) An architectural framework for media server control. Tech. rep., Internet Engineering Task Force, Request For Comments (RFC) 5567
48. Moody DL (2009) The physics of notations: toward a scientific basis for constructing visual notations in software engineering. *Software Engineering, IEEE Transactions on* 35(6):756-779
49. NUBOMEDIA Consortium (2015) NUBOMEDIA: an elastic Platform as a Service (PaaS) cloud for interactive social multimedia. Available online: <http://www.nubomedia.eu>, ICT-2013.1.6 (GA-610579)
50. O'Doherty P, Ranganathan M (2003) JAIN SIP API Specification. Tech. rep., Java Community Process, Java Specification Request (JSR) 32
51. Ott J, Wenger S, Sato N, Burmeister C, Rey J (2006) Extended RTP profile for real-time transport control protocol (RTCP)-based feedback (RTP/AVPF). Tech. rep., Internet Engineering Task Force, Request For Comments (RFC) 4585
52. Park S, Park NS, Kim JT, Paik EH (2009) Provision of the expressive multisensory adaptation platform for heterogeneous multimedia devices in the ubiquitous home. *Consumer Electronics, IEEE Transactions on* 55(1):126-131
53. Perkins C, Westerlund M (2010) Multiplexing RTP data and control packets on a single port. Tech. rep., Internet Engineering Task Force, Request For Comments (RFC) 5761
54. Piccioni M, Furia C, Meyer B, et al (2013) An empirical study of API usability. In: *Empirical Software Engineering and Measurement, 2013 ACM/IEEE International Symposium on*, IEEE, pp 5-14
55. Prusty N (2015) *Learning ECMAScript 6*. Packt Publishing Ltd
56. Reddy M (2011) *API Design for C++*. Elsevier
57. Rosenberg J (2010) Interactive connectivity establishment (ICE): A protocol for network address translator (NAT) traversal for offer/answer protocols. Tech. rep., Internet Engineering Task Force, Request For Comments (RFC) 5245
58. Rosenberg J, Schulzrinne H, Camarillo G, Johnston A, Peterson J, Sparks R, Handley M, Schooler E (2002) SIP: session initiation protocol. Tech. rep., Internet Engineering Task Force, Request For Comments (RFC) 3261
59. Saint-Andre P (2011) Extensible messaging and presence protocol (XMPP): Core. Tech. rep., Internet Engineering Task Force, Request For Comments (RFC) 6120
60. Saleem A, Xin Y, Sharratt G (2010) Media server markup language (MSML). Tech. rep., Internet Engineering Task Force, Request For Comments (RFC) 5707
61. Schulzrinne H (1998) Real time streaming protocol (RTSP). Tech. rep., Internet Engineering Task Force, Request For Comments (RFC) 2326
62. Schulzrinne H (2003) RTP profile for audio and video conferences with minimal control. Tech. rep., Internet Engineering Task Force, Request For Comments (RFC) 3551
63. Schulzrinne H, Rosenberg J (1999) The IETF internet telephony architecture and protocols. *Network, IEEE* 13(3):18-23

64. Schulzrinne H, Casner S, Frederick R, Jacobson V (2003) RTP: A transport protocol for realtime applications. Tech. rep., Internet Engineering Task Force, Request For Comments (RFC) 3550
65. Tang X, Zhang F, Chanson ST (2002) Streaming media caching algorithms for transcoding proxies. In: Parallel Processing, 2002. Proceedings. International Conference on, IEEE, pp 287-295
66. Taylor T (2000) Megaco/H.248: A new standard for media gateway control. Communications Magazine, IEEE 38(10):124-132
67. Thom G, et al (1996) H.323: the multimedia communications standard for local area networks. Communications Magazine, IEEE 34(12):52-56
68. Tilkov S, Vinoski S (2010) Node.js: Using javascript to build high-performance network programs. Internet Computing, IEEE 14(6):80-83, DOI 10.1109/MIC.2010.145
69. Van Dyke J, Burger E, Spitzer A (2007) Media Server Control Markup Language (MSCML) and Protocol. Tech. rep., Internet Engineering Task Force, Request For Comments (RFC) 4722
70. Vaquero LM, Rodero-Merino L, Caceres J, Lindner M (2008) A break in the clouds: towards a cloud definition. ACM SIGCOMM Computer Communication Review 39(1):50-55
71. Vogel A, Kerherve B, von Bochmann G, Gecsei J (1995) Distributed multimedia and QoS: a survey. MultiMedia, IEEE 2(2):10-19, DOI 10.1109/93.388195
72. Wagner B (2010) Effective C# (Covers C# 4.0): 50 Specific Ways to Improve Your C#. Pearson Education
73. Westerlund M, Wenger S (2016) RTP internet draft topologies draft-ietf-avtcortp-topologies-update-10. <https://tools.ietf.org/html/draft-ietf-avtcortp-topologies-update-10>, accessed 7 January 2016
74. Westerlund M, Burman B, Nandakumar S (2015) Using Simulcast in RTP Sessions draft-westerlund-avtcortp-simulcast-04. <https://tools.ietf.org/html/draftwesterlund-avtcortp-simulcast-04>, accessed 12 December 2015
75. Willmott S, Balas G (2013) Winning in the API Economy. Available online: <http://www.3scale.net/wp-content/uploads/2013/10/Winning-in-the-API-EconomyeBook-3scale.pdf>
76. World Wide Web Consortium (2011) Web real-time communications working group. <http://www.w3.org/2011/04/webrtc/>, accessed 11 December 2015
77. World Wide Web Consortium (2015) WebRTC Stats. Available online: <http://w3c.github.io/webrtc-stats/>
78. Zhou F, Duh HBL, Billingham M (2008) Trends in augmented reality tracking, interaction and display: A review of ten years of ISMAR. In: Proceedings of the 7th IEEE/ACM International Symposium on Mixed and Augmented Reality, IEEE Computer Society, pp 193-202
79. Zhu W, Luo C, Wang J, Li S (2011) Multimedia cloud computing. Signal Processing Magazine, IEEE 28(3):59-69