

Networked Multimedia with Internet Media Guides

Jörg Ott¹, Dirk Kutscher², Janico Greifenberg², Kevin Loos²

¹*Helsinki University of Technology, Networking Laboratory
e-mail: jo@netlab.hut.fi*

²*Technologiezentrum Informatik (TZI) Universität Bremen
e-mail: {dku|jgre|logic}@tzi.uni-bremen.de*

Abstract: Internet Media Guides (IMGs) are a generalisation of Electronic Programme Guides (EPGs) as known from digital video broadcasting (DVB). They are independent of specific metadata formats and thus are able to support a broad range of applications, including EPG distribution for TV networks and distribution of session descriptions for Internet-based multimedia sessions. Unlike most existing approaches, the IMG framework is also completely independent of specific delivery networks for the media content described in media guides—and it is also independent of the distribution mechanisms for the media guides themselves: IMGs can be distributed in unidirectional broadcast networks, they can also be retrieved over established query/response protocols such as HTTP, and they allow for asynchronous change-notifications to interested subscribers. This paper presents the fundamental concepts of the IMG framework and introduces the main components. We further describe implementation and application experiences and discuss potential enhancements to the existing IMG specifications.

Key words: Internet Media Guides, multimedia communications, session description

1. INTRODUCTION

Research networks have played a major role in developing and advancing internetworked multimedia technologies in the past: Early real-time multiparty, multimedia conferences have been performed between universities, based on Mbone multicast connectivity, research projects have developed and validated architectures and applications for multimedia conferencing, tele-learning, and IP telephony. Later, practical deployments of Internet multimedia built upon these experiences, but often not exploiting the full potential of the technologies. While IP telephony has been flourishing for quite some time, media dissemination has largely been constrained to (HTTP-based) streaming linked from web pages and Internet radio sites. SDR-style [Han96] directories of streaming offerings did not really exist; services usually relied upon human-readable information embedded in web pages.

Only recently, with the advent of digital mobile broadcasting technologies (DVB-T/H), broadcasting services for 3G networks, the introduction of video (on-demand) services for residential users, and the evolution of networked multimedia home architectures, this landscape has started to change. The largely independent deployment of these – often, but not always IP-based – delivery methods has led to a split into different application areas, each of which relies on dedicated networks, transport and control protocols, and finally, different approaches to describing media sessions and to distributing these descriptions to user devices. Clearly, the assumption of a single, unified IP multicast-enabled architecture, that the original SAP/SDP- and RTP- based approach

to media broadcast is based on, is no longer valid. Instead, new approaches must be developed that can achieve an integration of service- and content-access across different networking environments.

The Internet Media Guide (IMG) framework, a current development and standardisation effort in the IETF, is an important component in that respect and may be used as replacement for SAP/SDP for media broadcast applications but goes well beyond the scope of SAP/SDP. The IMG framework supports common description languages to express meta information about services (and contents), their identification and the dissemination of such information in a standardised fashion. Peered with adequate description languages, Internet Media Guides (IMGs) [NWL⁺04b] [NWL⁺04a], are able to provide the necessary service description and discovery functionality that is required for today's diverse media broadcast environment and will also support the incremental conversion to an integrated IP-based infrastructure.

This paper describes the IMG framework and highlights the potential and the challenges (not just) for National Research and Education Networks (NRENs) and the networking research community when upgrading from traditional Internet Multimedia Conferencing Architecture-based media broadcast (IMCA) [HCBO99] to a generalised, integrated media broadcast environment. While the fundamentals of this architecture remain valid, two relevant changes can be identified – we address their implications on session description mechanisms in this paper: 1) No Internet-wide multicasting in the sense of Any Source Multicasting will widely be available. Rather, a combination of point-to-point communications (possibly

organised as application-layer multicast overlay) and Source-Specific Multicasting (SSM) [HC04] will be dominant, possibly peered with link/sub network-local multicast/broadcast. These communication mechanisms will be used for disseminating control as well as media traffic. 2) No exclusive reliance on IP addresses will be possible as long as different link layer technologies (still) require network-specific addresses.

This paper is organised as follows: Section 2 introduces the IMG framework in more detail and points to related work where appropriate. Section 3 describes our implementation of a generic IMG platform and Section 4 shows a sample application we have built on top of this platform. Finally, Section 5 concludes this paper with a brief summary and hints at future work.

2. CONTENT AND SERVICE DESCRIPTION WITH IMGs

Internet Media Guides (IMGs)¹ are a generalisation of Electronic Programme Guides (EPGs) as known from digital video broadcasting (DVB) applications such as TiVo², and the TV Anytime forum [The05]. The IMG generalises EPGs in at least three ways:

- No restrictions to (TV) programme schedules – IMGs support arbitrary description formats so that other descriptions of other media types (including static web pages) and services may be carried in IMGs as well.
- Independence of delivery networks – IMGs as well as the content may be distributed across a variety of different networks (while preserving the content identification). Where necessary, media descriptions are expected to provide content locators for multiple delivery networks.
- Flexible distribution mechanisms – the IMG framework supports a set of complementary delivery mechanisms capable

of serving users on one-way and two-way networks, workable over a wide range of data rates, and providing asynchronous dissemination and notification as well as on-demand retrieval.

The IMG framework [NWL⁺04a] builds upon well-established Internet technologies and restricts its added features to a minimum set to achieve the aforementioned properties. These additions comprise support for arbitrary content metadata and a set of mechanisms for IMG transport. The left-hand side of Fig. 1 gives an overview of the IMG framework and its constituents: Metadata describing contents and services – which make up IMGs – are made available in arbitrary formats and encapsulated in an IMG envelope for dissemination. IMGs originate from IMG senders and are inquired for and/or conveyed by means of IMG operations using a set of suitable transports to IMG receivers. Optional intermediaries – IMG transceivers – may combine, manipulate, filter, personalise, and otherwise process IMGs on their way. Security mechanisms are developed to ensure that, *e.g.*, integrity and authenticity of IMGs may be preserved in spite of intermediaries.

The following three subsections address the envelope and identification concept and the IMG transport operations,

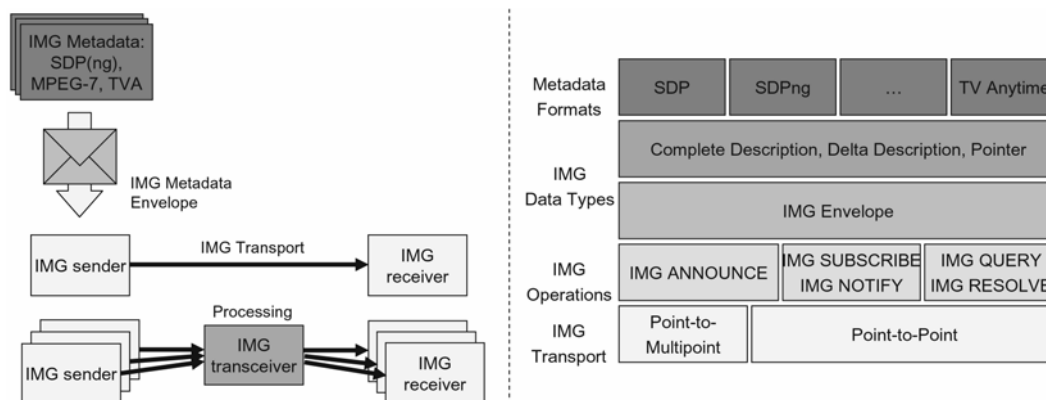


Fig. 1. Architecture overview

respectively, and outline how cross-network integration can be achieved using IMGs.

2.1. Supporting application-independent descriptions

IMGs are independent of the actual metadata contents they carry because suitable service-descriptions depend on the actual application and cannot be easily generalised due to the diversity of applications. Examples include: broadcast scheduling (*e.g.*, TV Anytime), old Mbone sessions (SDP [HJ98], SDPng [KOB05]), web resources (plain HTTP and RTSP URLs), WLAN access (WiFi, IEEE 802.21), triple play services offerings, among others. Therefore, the IMG approach restricts itself to: 1) defining three elementary data types workable with all metadata formats; 2) providing an envelope format with minimal functionality to allow content-

¹ The specifications are currently being developed in the IETF MMUSIC WG in coordination with 3GPP and DVB; their completion is anticipated for late 2006.

² <http://www.tivo.com>

-independent processing on the way and to provide the required security properties; and 3) an identification scheme for IMGs based upon URNs.

IMGs are deliberately defined in an open sense in how much information they may comprise: an IMG may contain as little as a service announcement of a hotspot or single SDP-based session description of a few hundred bytes but may as well contain a full month of programming information of a cable TV station serving 100 channels including background information amounting to several megabytes of metadata. To provide flexibility when distributing and updating information within this broad range, three data types are defined for IMGs: Full IMGs contain a complete set of metadata from the perspective of the respective originator and are expected to be distributed regularly or at least available upon request. Delta IMGs comprise the difference with respect to a reference Full IMG and are used for IMG updating. Finally, IMG pointers are references to IMGs or IMG fragments embedded elsewhere in a message or stored in an external location. They may be used, *e.g.*, when informing receivers about updates asking them to fetch the updated information whenever convenient.

The IMG envelope [WLP⁺05] is used to encapsulate arbitrary metadata formats in one of the three representations described above. The IMG envelope allows denoting the content and data type, provides versioning and validity (particularly expiration) information, and supports unique identification of the contained metadata. Specific attributes also identify the delta encoding used (if any) and may indicate fragments of a Full or Delta IMG. The IMG envelope also supports security properties: envelopes may be signed and/or encrypted as a whole and, if supported by the envelope format, in part.

Currently, two alternative approaches are discussed to implementing IMG envelopes: Multipurpose Internet Mail Extensions (MIME) and XML. XML-based envelopes allow for a potentially more fine-grained encapsulation and identification, particularly when the metadata formats themselves are XML-based; they may draw upon specifications for, *e.g.*, XML signatures to provide security features. However, this approach is less efficient with binary metadata as the metadata would require extra encoding. In contrast, the MIME-based approach can carry arbitrary data and supports security by inheritance from S/MIME [Ed04]. However, overhead may increase substantially if fine granularity for embedding metadata is required. Note that none of the approaches can be defined completely based upon existing technologies: XML-based envelopes require a new schema to be defined while MIME envelopes need additional MIME headers and/or parameters. 3GPP has chosen to use XML but defines only a very basic set of IMG functions (without any security) for its MBMS [3GP04]; no decision has yet been taken in the IETF but both approaches appear reasonable and the present IETF envelope-specification is backwards-compatible to 3GPP but leaves room for other envelope-encodings than XML.

IMG envelopes are presently identified by the MIME type “application/envelope+xml” in case they follow the XML-

-based specification; all other envelope formats will be registered as “application/envelope.regid”, with “regid” being assigned in the registration process (*e.g.*, “mime” in case of a MIME envelope). The precise approach is still under discussion at the time of writing [WLP⁺05].

Finally, IMG instances need to be identified: those sourced by the origin servers as well as those created by transceivers by processing incoming IMGs (and combining, splitting, or otherwise modifying them). Unique Resource Names (URNs) [Moa97] are chosen for this purpose and mechanisms will be provided to specify versioning information in a standardised manner. Any intermediate operation that does not preserve the structural integrity of an IMG (*i.e.*, combines or splits IMGs rather than just modifying fields) creates a new IMG-instance with a new URN. URNs may be peered with location-information – about (the domain of) the entity that created it – to enable its later retrieval without requiring a separate URN resolution infrastructure.

Delta encoding and IMG fragments pose particular requirements on the identification and encapsulation of IMGs. The construction of deltas is specific to the respective metadata format and needs to be specified separately – and the respective delta encoding method used must be explicitly indicated in the envelope. So far, no standard delta mechanisms have been specified. We have experimented with methods and formats operating on line-based text formats [Gre05]: following the Unix tools `diff` and `patch`; with algorithms [CRGMW96] and encodings such as `Xupdate`³ specifically tailored to structured XML-based description documents; and with tools such as `xdelta`⁴ – a generic delta encoding for arbitrary (binary) description formats. The identification of metadata fragments also depends on the metadata format in use: if an XML-based format is used, XPath [CD99] may be used to identify the location within the XML document tree where the respective delta is to be applied. This also allows provision of an increasing level of detail over time about, *e.g.*, a service or programme schedule. We have chosen to make use of a generic anchor format for fragment identification that may be populated with XPath information (and is interpreted as such if the metadata format so indicates).

IMGs may be modified by transceivers by adding/removing, combining, or altering parts. To maintain integrity and authenticity properties in spite of transceiver operation along the way from an origin IMG sender to the ultimate receiver, a modification history may be attached to IMGs, which allows a receiver to trace the announcement back to the origin and validate this IMG’s integrity for each step. Such modifications may be explicitly included as delta-encoded IMG fragments or (pointers to) the original version may be included along with the respective signatures. Further details will be investigated as the IMG envelope specification proceeds.

³ <http://xmldb-org.sourceforge.net/xupdate/>

⁴ <http://xdelta.org/>

2.2. IMG access and distribution

The metadata-related aspects described in the previous section are independent of the underlying transport mechanisms as depicted in the IMG architecture in Fig. 1 on the right-hand side. Three basic IMG operations are defined to address access to and distribution of IMGs in a scalable fashion and serve large numbers of a variety of receivers in diverse networks.

IMG ANNOUNCE is used for IP-multicast-based dissemination of IMGs. To be able to work with large messages, with large receiver groups, and in unidirectional networks, IMG ANNOUNCE is based upon the reliable multicast file delivery protocol FLUTE (RFC 3926) which, in turn, makes use of the reliable multicast protocol for an asynchronous layered coding (ALC, RFC 3450) that uses the layered coding transport (LCT, RFC 3451) peered with forward error correction (RFC 3452, RFC3453).

ALC defines (inherited from LCT) the notion of distribution channels (which coincide with multicast transport addresses) across which data objects are conveyed. A set of such channels paired with a sender identification makes up an ALC session. FLUTE (RFC 3926) builds upon this mechanism to form a FLUTE session. FLUTE defines the transmitted objects to be files, with one dedicated object per FLUTE session, the File Description Table (FDT), conveying an overview of all the files transmitted along with characteristic information such as filename, file size, last modification time, among others, as a simple XML document. The MUPPET specification [LPP03] adopts FLUTE for IMG transmission: The IMG ANNOUNCE sessions can be described in SDP [HJ98] as specified in [WCPP05] thus enabling system bootstrapping but also cross-references between different IMG ANNOUNCE sessions.

Within MUPPET, different distribution channels are used for Full and Delta IMGs embedded in envelopes. An IMG Announcer transmits IMGs as identifiable objects within a FLUTE session in sender-controller intervals using a carousel-style mechanism (which is up to the discretion of the announcer) and may spread out the transmission across multiple layers to allow for receiver-controlled congestion control. IMG Listeners join a predefined (or otherwise obtained) IP multicast group per FLUTE session and determine which IMG objects are transmitted on which multicast addresses. The listeners accumulate the received information and regenerate the full IMG representation. Appropriate URIs referring to a FLUTE session still need to be defined.

In case IMG fragments are missing (*e.g.*, because the listener was turned off or because of packet losses) or if no multicast channel is available at all, a complete IMG or parts of it need to be retrieved interactively. The IMG receiver may fetch a (fragment of a) Full or Delta IMG. The IMG QUERY operation and corresponding IMG RESOLVE response are mapped onto an HTTP GET request-response pair [OL05]. The querier issues an HTTP GET request specifying location

(an HTTP URL correlated to the URN sought) and provides additional information (*e.g.*, version numbers) encoded in the URI or its parameters. Regular HTTP headers may be used to specify subsets of data to be retrieved (Range:), to perform conditional fetching (If-Modified-Since:), among others. The IMG resolver returns the specified IMG embedded in an envelope in the 200 OK response (or an appropriate error code). Other protocol mechanisms applicable to HTTP, *e.g.*, using TLS as secure transport or dynamic redirection are directly applicable as well.

Finally, IMG receivers may desire asynchronous notifications when IMGs change – rather than having to poll regularly or wait for a carousel transmission. For this purpose, the IMG framework supports IMG SUBSCRIBE and NOTIFY operations that may be mapped to the corresponding SIP operations (RFC 3261, RFC 3265) with appropriate event packages yet to be defined [NS04]. Other notification mechanisms are conceivable as well (we have carried out some experiments with Jabber as notification infrastructure). Constructing SIP URIs for subscriptions and the definition of filtering rules for SIP SUBSCRIBE message bodies appears straightforward: the domain name of the subscription SIP URI is taken from the originator of the IMG to be subscribed, the username portion comes from the URN. This will result in all changes being reported and all versions to be made available. The subscriber may use elements in a yet-to-be-defined subscription message body to limit subscriptions: to subsets of the IMG and/or to certain versions, intervals, validity times, etc. The subscriber should be able to identify whether to receive full or delta IMGs or just IMG pointers. Given the only moderately dynamic nature of IMGs, we assume that event throttling will not be a major issue. Finally, an optimisation may allow a subscriber to specify whether it would be capable of receiving IMG notifications in a FLUTE session and, if so, allow the notifier to provide a pointer to such a session; this would allow efficient distribution of mass notifications at the discretion of the notifier.

Figure 2 shows how an IMG delivery system may be set up: two IMG senders gather metadata from various content providers and offer them via different distribution channels. The IMGs are aggregated at a content broker that forms a coherent programme offering for a certain distribution network. A regional provider (*e.g.*, in a county) adds local programming information to the bouquet. As access service provider, it allows its customer to provide “filter” specifications, so that the vast content choice is reduced to meet the respective customer’s needs. While the information distribution up to the service provider was performed *via* multicast distribution, the last hop to the customer uses individual delivery mechanisms, at least for the personalised parts. The receiver may have it’s the personal IMG profile agreed upon with the service provider. At power-on, the IMG receiver fetches the full personalised IMG and then subscribes to it to be notified about changes. Alternatively, the receiver may also listen to IMG multicast of certain genres, perform

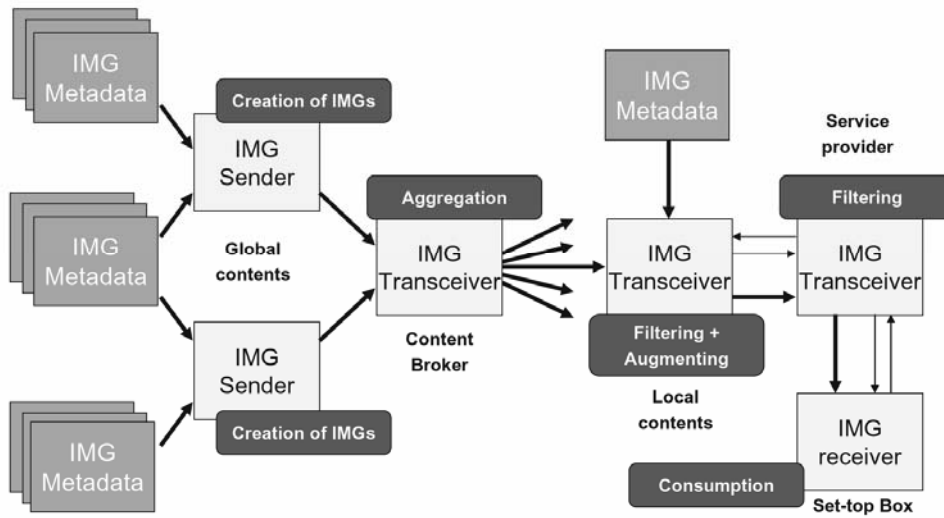


Fig. 2. Sample IMG operation with metadata aggregation and personalisation

filtering locally, and initiate subscriptions with the service provider only for a few crucial pieces of metadata.

2.3. Network infrastructure integration with IMGs

As mentioned in the introduction, IMGs may be used to bridge different delivery networks. While Figure 2 sketches the basic operation of an IMG system, it does not address this

integration aspect. The IMG receiver in the above figure, however, may well be a portable device capable of accessing multiple networks: for example, a small TV set supporting DVB-T and WLAN. The service provider for IMG delivery may have access to IP streaming content providers as well as the regional DVB operators; there may even be multiple service providers offering the same content. With DVB-T

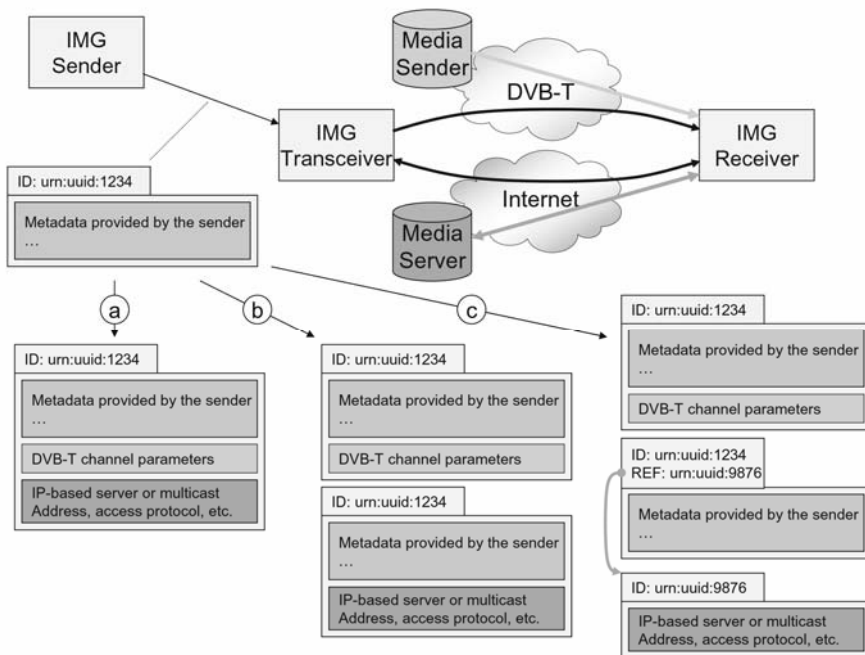


Fig. 3. Using IMGs to integrate access to content via different service providers

presently using classical MPEG streaming for video broadcasts, a content description, for example, of the tagesschau (which is probably the most popular German news magazine on TV) needs to be created only once (the “session” level in SDP) but different transports would be specified (the “media” level in SDP). For example, the pairs of DVB-T regions and the respective channel/multiplexing parameters may need to be provided for DVB-T as well as the IP addresses for three different data rates plus a layered video coding session for source-specific multicast. The transport addresses of the streaming servers may even be adapted by regional service providers to direct clients to the closest media caches; while providing multiple server addresses for load balancing is conceivable at the IMG layer, this may as well be left to the DNS or schemes embedded into the control protocol (*e.g.*, the redirect function of HTTP or RTSP).

Figure 3 shows a conceptual outline of such a set of IMG metadata and its modification in a transceiver to fit two different distribution paths. If a metadata format is capable of expressing several alternative ways of accessing media content (as is the case with SDP and SDPng) the IMG framework restricts its support to encapsulation and support for modifying or augmenting the metadata along the way (Fig. 3a). If a metadata format does not support, for example, DVB-T reception parameters as well as IP addresses at the same time, the IMG envelope concept may be used to provide identifiable modifications of the original description, namely substitution of one transport address by another – which may be encoded as a delta representation (Fig. 3b). Finally, if the metadata format does not support all the necessary address types, IMG extensions are conceivable (but not yet specified) that allow linking (a fragment of) the metadata in one format to supplemental information represented in another (Fig. 3c). In all cases, the IMG framework allows metadata to be retrieved *via* any of the involved networks and provides appropriate delivery mechanisms for all of them.

3. IMPLEMENTATION

We have implemented a platform for IP-based distribution of Internet Media Guides that can be employed for different application scenarios⁵. The platform supports creation and management, distribution and retrieval as well as manipulation and forwarding of arbitrary IMGs. The implementation is split into two parts: a content-independent distribution platform called Papageno⁶ [Loo05] and an IMG management system that uses Papageno as IMG transport.

Papageno has been developed as a general-purpose distribution and reception subsystem that can be linked to application-specific programmes. Currently, it supports MUPPET/-

FLUTE-based distribution of arbitrary objects, including IMG delta descriptions and pointers. Additional mechanisms are work in progress: an HTTP mapping for IMG QUERY/RESOLVE has been developed but is not yet integrated; IMG SUBSCRIBE and NOTIFY have been tested using Jabber; the transition to SIP is in progress.

Papageno encompasses both sender and receiver side for all IMG operations, and it supports transceiver operation. The FLUTE sender can be used for distribution of arbitrary objects, *e.g.*, large files containing IMG metadata, and can be configured to automatically perform delta computation for different versions of the same object. When an application instantiates a FLUTE session it can specify different parameters such as the number of complete channels and delta channels, rate limits for channels, and the FEC encoding to be used. The sender can then generate an SDP description (using the format specified in [WCPP05]) for the FLUTE session reflecting these parameters. This SDP description can be made available to receivers, *e.g.*, by distributing it out-of-band for bootstrapping or as part of IMG metadata (transparent to Papageno) for cross-referencing between IMG distribution channels represented by FLUTE sessions.

The IMG management system is also implemented as a library to be used by various applications. The IMG library generates and processes the IMG envelope and handles delta IMG descriptions. Another important function is the storage and management of the available metadata. This includes the current versions as well as older ones needed for creation and processing of deltas. If the store is missing a certain piece of metadata, *e.g.*, if the old version needed to process a delta IMG description was not received so far, the store can explicitly request Papageno to retrieve the metadata.

The store component is used in both senders and receivers: On the sender side, the application adds new metadata to the store which passes the metadata to Papageno. On the receiver side, Papageno inserts received metadata into the store and the latter delivers the metadata to the application. In case of a transceiver, Papageno is used for both sending and receiving but there is only one store. The application part in a transceiver modifies the metadata (*i.e.* performs any of the transceiver operations mentioned above such as aggregation, filtering, etc.) after it has been received and before it is sent.

Figure 4 illustrates how the IMG platform performs the IMG operations. The black lines represent the messages passed in a sender-driven distribution of the metadata (push-distribution); the blue lines represent the exchange in case of a receiver-driven explicit fetching (pull-distribution).

In both cases, the initial step is identical: creating or updating the metadata. The application passes the metadata to the store component of the IMG library which inserts the metadata into its internal data structures (Step 1). The store informs the transport implementation Papageno about the availability of new metadata and passes the metadata as complete description, delta description, and pointer (Step 2).

⁵ More information about this implementation is available at <http://www.dmn.tzi.org/research/img/>

⁶ The name has been motivated by FLUTE as the primary transport implemented.




Fig. 4. Overview of IMG platform

The MUPPET implementation in Papageno performs the IMG ANNOUNCE operation sending the three IMG data types through different FLUTE channels. The Papageno notification mechanism sends an IMG pointer to each subscriber, *i.e.*, the receivers who sent an IMG SUBSCRIBE message to the sender (Step 3).

When an IMG receiver wants to explicitly fetch metadata it sends an IMG QUERY message, *i.e.* an HTTP GET request (Step 2). In the IMG platform, the Papageno HTTP implementation is responsible for handling this kind of request. Papageno queries the store component in the IMG library (Step 3) that returns the metadata with an IMG envelope, provided the metadata has been added by the application (Step 4). Metadata and envelope are sent to the querier in the body of the HTTP response (Step 5).

On the receiver side, the IMG platform components are used in reverse order: Papageno receives the metadata from the network and adds it to the store in the IMG library. This component checks whether the metadata it really a new version (and ignores duplicates) and passes new metadata to the application.

4. APPLICATIONS

Leveraging our IMG implementation, we have applied IMGs to several application scenarios: distribution of SAP sessions announcements, distribution of electronic programme guide information for TV programmes, and distribution of network service descriptions, *e.g.*, for supporting network access in public WLAN hot-spots.

The SDP/SAP framework for distributing information about multicast multimedia sessions in the Internet is still being used for announcing sessions (and for co-ordinating multicast address usage). SAP session announcements are periodically sent to a standardised IP multicast address and their distribution can be restricted to scopes (*via* IP TTL scoping or administratively configured IP multicast scopes). By relying on IP multicast as the sole distribution mechanism, SAP is not useful when IP multicast and Mbone connectivity is not available (particularly if access to the advertised contents is possible *via* unicast, as can increasingly be observed today).

In order to make SAP announcement accessible to a broader audience, we have developed a SDP/SAP-IMG gateway that receives SAP announcements and can distribute them over the different IMG distribution mechanisms (IMG operations). For example, users can join the multicast session for MUPPET/FLUTE-based distribution, they can retrieve current announcements *via* HTTP, and they can subscribe for SAP/SDP announcement- and receive-notifications about changes.

Our current implementation translates one session description that has been received *via* SAP to a single media guide that is made available *via* IMG operations. For IMG ANNOUNCE distribution using MUPPET/FLUTE, we use a single MUPPET session that is used to convey the different media guides (as well as deltas and pointers).

While the SDP/SAP gateway is useful to make session announcements more accessible, *e.g.*, to home users lacking multicast connectivity, it does not solve the problem of providing access to the actual content that is announced, *i.e.*, the multicast media sessions. Providing access to multicast sessions for users with only unicast connectivity, it is typically implemented by multicast/unicast gateways that forward multicast packets to the unicast addresses of specific subscribers. Often, this function is coupled to a transcoding gateway that can also transcode the media content, *e.g.*, in order to accommodate users with limited device capabilities or network resources.

In fact, multicast/unicast (and transcoding) gateways fit quite well to the IMG transceiver concept, as a media gateway has to inform potential receivers about available media sessions. By co-locating an SDP/SAP-to-IMG gateway and a media gateway, the IMG distribution function can act as an IMG transceiver and augment the original session description with specific transport parameters for accessing sessions over the media gateway. Such a system could also offer different variants of a single original session, *i.e.*, different transcoding parameters or different transport mechanisms.

We are currently extending our IMG-based SDP/SAP gateway to a complete media relay. First experiences with the stand-alone SDP/SAP-to-IMG function have shown that mapping each received session description to an independent media guide is probably the right decision for providing

a SAP replacement because this approach would help applications such as session directory tools to incrementally update their list of sessions in an easy way. However, for a general, large-scale approach, some form of grouping (*e.g.*, into genres) could be useful. In addition, for some IMG application scenarios, *e.g.*, users that access IMGs *via* IMG SUBSCRIBE and IMG QUERY, one might want to provide a filtering mechanism that allows users to subscribe to update notifications for only a subset of the available media guides. Furthermore, we have noticed that, today, SAP updates are not really common. Instead, session announcements are typically created by an announcer and in case changes occur, the original announcement is simply deleted (or times out), and a new announcement is created, which makes it difficult for an IMG gateway to identify updates and to make use of the delta mechanism for updates.

IMGs are particularly interesting for broadcast applications in networked multimedia scenarios. For example, digital personal video recorders (PVRs) allow for recording selected content (from TV networks) for later viewing. One of the popular functions of PVRs is to automatically select content based on user preferences, *e.g.*, by evaluating filter expressions. In order to identify and select content, PVRs need programme guide information, *e.g.*, the title of individual programmes, scheduling information, and meta-information about the programme.

Different broadcast networks provide different support for distributing programme guides for the aired content, *e.g.*, DVB provides periodic distribution of programme guide information for current and future programmes, whereas analogue broadcast networks provide no or only a limited form of programme guide information. In order to provide PVRs with meaningful programme guide information reliably, it can therefore be useful to leverage other, TV-broadcast-independent distribution mechanisms, *e.g.*, programme guide distribution over the Internet, as employed by the commercial TiVo product.

Whereas proprietary solutions can solve the programme guide distribution problem for specific TV networks (and for specific markets, *e.g.*, the US in the TiVo case), a more generalised approach for network-independent programme guide distribution allows to support a broader range of content providers, *i.e.*, not only classic TV broadcast networks, and thus provides a broader selection of content for users.

Based on the IMG distribution platform, we have built a programme guide distribution application that can be used to provide applications similar to TiVo with programme guide information. The general operation is depicted in Fig. 5: programme guide information is extracted from public sources, *e.g.*, web sites, and converted to a canonical format. The current implementation uses the XMLTV format⁷ – a simple XML-based format for describing programme information. The generated XMLTV programme guide is then provided as input to the IMG distribution platform where it is encapsu-

lated in IMG envelopes and made available over the different IMG distribution mechanisms (operations). For example, clients such as users' PVRs can subscribe to selected programme guide information and then receive updates (or notifications about updates). As a sample application, we have used Freevo⁸ – an Open Source front end for home theatre applications. Freevo can present programme guide information to users and process it, *e.g.*, in order to provide time-shifting services.

There are different possibilities for structuring EPG information for IMG-based distribution, *e.g.*, one could aggregate all available programme guides into a single IMG document and update this media guide as soon as additions or changes are available, or one could define a more fine-grained structure and generate individual media guides for individual programmes or genres. The design of a concrete IMG distribution system largely depends on the specific application requirements, *e.g.*, media guides for presentation to humans provide different requirements than media guides for automatic processing by PVRs. In our implementation, the main objective was to provide Freevo with complete programme guide information, in an efficient way. Therefore, we have decided to generate a single media guide for all programmes. It should be noted that, given an appropriate metadata format, it is always possible to employ IMG transceivers that take a complete media guide and then generate customised media guides, *e.g.*, by splitting media guides into individual guides per programme or genre.

TV programme guides can be obtained from different sources, *e.g.*, from TV network or TV magazine websites, from EPG feeds within the TV network (for DVB networks), or they can be edited manually. From our first experiments with IMG-based distribution of TV programme guides, we have observed that, today, EPGs are largely static content and do not change after they have been created and published (this holds for both, EPG information on websites and for DVB-EPGs). Late-time changes to the programme schedule are often not reflected in any EPG, not even in the online EPG of the DVB network. As a result, changes to the IMG typically occur only as additions, not as updates to existing media guides. This may be related to the high cost that is currently associated with maintaining and delivering programme guides over different networks with different distribution mechanisms without a unified distribution and processing architecture. The introduction of IMGs for this application would make it easier to deal with dynamically changing media guides and to deal with updates more efficiently, which could potentially change the way media guides are handled both by producers (content providers) and consumers (end systems).

There are some interesting practical considerations when deploying IMGs for real applications. For example, how do IMG receivers learn about available IMGs and potential ways to access them? For SAP-based session-announcements on

⁷<http://membled.com/work/apps/xmltv/>

⁸<http://www.freevo.org/>

theMbone, this question can be easily answered: There is one well-known multicast group (respectively: one multicast group per administrative scope) that receivers can tune into in

desirable to associate locators with these names. For example, the latest IMG for the radio channel “BBC1” might be associated with a fixed, well-known URN, and in addition,

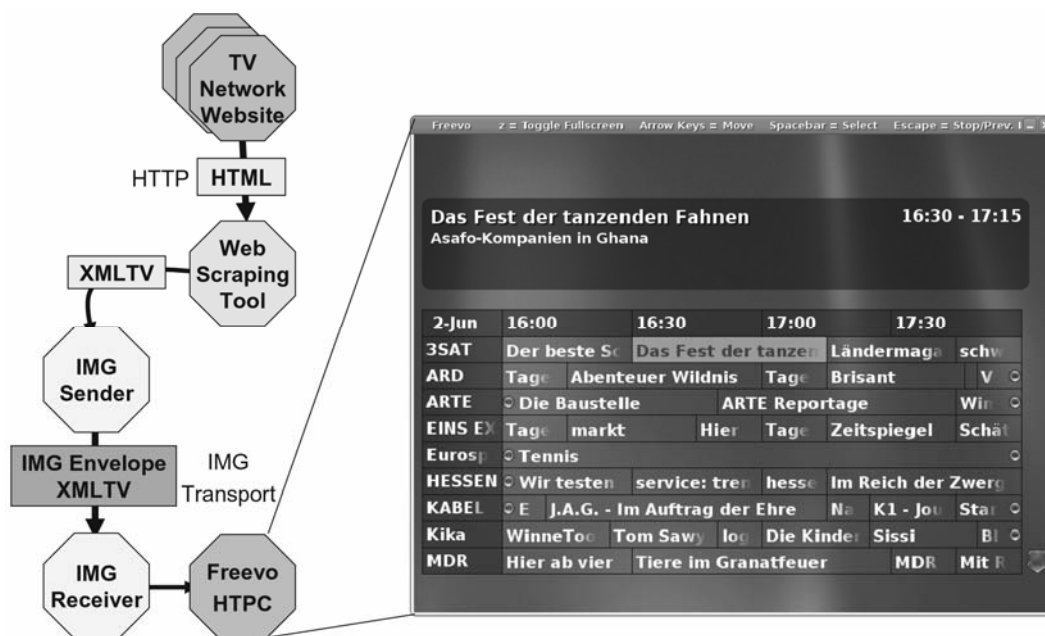


Fig. 5. TV-EPG distribution

order to receive multicast session announcements. Clearly, for IMGs as a generalised platform for metadata distribution, such an approach is neither viable nor particularly desirable, since IMGs are not limited to a single network, *e.g.*, the global IPv4-based Internet, but are intended as a network- and distribution mechanisms-independent technology that can be employed for a broad range of applications. For some applications, *e.g.*, for IMG-based distribution of TV programme guides in the Internet, an IMG sender application could be conceived as a server that acts as a service discovery tool and can be accessed using a well known URI. For example, an HTTP server could provide a directory for the available IMG services and link to session descriptions that allow receivers to join the MUPPET/FLUTE sessions. For broadcast-only environments, *e.g.*, IMG distribution over DVB data services, an explicit query-response scheme for IMG discovery is not viable; instead the bootstrapping would have to be implemented using a broadcast/multicast mechanism. For example, one could imagine a dedicated (and well-known) service discovery channel that is used to transmit meta information about available IMG channels, *i.e.*, IMGs about IMGs.

Independent of the concrete design decision that is taken for a specific application, this discussion shows that there is a need to better support IMG service discovery and to be able to describe how an specific IMG can be accessed. The existing IMG framework already provides the notion of unique IMG identifiers. In addition to that, it might be

one might want to specify one or multiple locators (each could be valid in a different network, *e.g.*, Internet or a digital radio broadcast network). There are some issues that have to be considered in this context, *e.g.*, the number of parameters for a MUPPET session is too high to express them in a single URI; instead one would rather have a URL that points to a description document. There is already a specification for expressing parameters for a single FLUTE session with SDP [WCPP05]; however, a comprehensive session description for IMGs is yet to be developed.

5. CONCLUSIONS AND NEXT STEPS

The IMG framework is a cross-platform service enabling technology for media broadcast in today's heterogeneous distribution networks. It is general and extensible enough to be applied in different network and distribution scenarios, ranging from large-scale distribution to site-local dissemination and adaptation to personal entertainment organisation and can thus spark the development of new multimedia broadcast applications and distribution frameworks.

IMGs' generality stems from three design characteristics: 1) independence of specific metadata formats, 2) independence of the actual content distribution, and 3) independence of specific networks and distribution mechanisms for the media guides themselves by providing different IMG operations. In addition, the concepts of IMG fragments, delta

IMGs, and IMG pointers allow for accommodating a broad range of applications, *e.g.*, from SAP-style periodic broadcast to decomposed programme guides with incremental updates, distributed in bandwidth-constrained environments.

First applications, such as the described SDP/SAP-gateway or the application of IMG within 3GPP's MBMS, largely deploy IMG as a multicast distribution mechanism for one-to-many distribution of relatively simple session descriptions – which means that the true potential of IMGs has not yet been practically explored. In fact, the current specifications leave some interesting questions open. First of all, the envelope format is currently application-specific, as there is no single mandated format that must be supported by any implementation. This may impose problems for generic IMG components such as IMG transceivers that are intended to work application-independent. We expect domain-specific standards to fill in for the respective area of application – *e.g.*, TV Anytime specifications for programming – and evolve over time to realise the full potential of the technology platform and enable integration of, so far, disjoint areas such as MBMS and TV service offerings. And, ultimately, we expect such programme guides – which describe just a set of homogeneous service offerings – to also enable service and service provider discovery and selection as a natural enhancement of content-based choices: from the fixed Internet access provider, to the wireless Internet service provider, to the media content broker, and to the service and contents.

The discussion of bridging different delivery networks in Section 2.3 has highlighted a use case of IMGs where a media guide contains both meta-information (about the content) and configuration information (technical information required for accessing the content in specific networks) and the media guide is distributed over various networks for describing content that is available over different networks. Even if the metadata format was expressive enough to support the specification of access information for different networks, it is not guaranteed that the original author of the media guide can anticipate all (future) distribution channels and thus include all the required configuration information. Instead, there might be something like relays down the distribution path that independently decide to offer the content *via* some specific network, *e.g.*, a local multicast distribution of a unicast streaming session. In this case, a new version of the original media guide would have to be created – for which we have discussed different options that could be pursued.

Finally, the notion of IMG transceivers provides for IMG elements that can generate new media guides based on existing ones, *e.g.*, by aggregating a set of input IMGs into a single output IMG or by adding local information to an IMG that was intended for larger-scope distribution. One of the interesting questions in this context is: how can the authenticity of the original IMGs be maintained over one or multiple transceivers in an IMG distribution path? Again, proposed solutions exist, but still need to be evaluated in large-scale applications and to be standardised.

At the moment, IMGs are a promising technology with a base set of specifications but more practical experience, especially in non-trivial application environments, is required. Similar to the pioneering efforts with SDP/SAP- and RTP-based Mbone sessions in the early days of real-time multimedia communication over the Internet, the NRENs and the connected institutions could play an important role in providing an ideal development and experimentation environment for this technology. Similar to the media broadcast environment, the academic research networks have evolved into a more heterogeneous environment as well, now providing different types of networks, usage scenarios and operating requirements. Backbone networks are adopting native IPv6, campus networks provide increased complexity, *e.g.*, comprising fixed Ethernets, IEEE 802.11 WLANs, roaming infrastructures, *etc.* New requirements for mobility and security imply new consideration for scalable and high-quality media distribution as well.

The highly evolved European research networks can provide a comprehensive test- and development infrastructure that allows researchers to develop tomorrow's large-scale media broadcast applications – the Internet Media Guide can be viewed as a first building block enabling the transition from traditional ASM-based real-time multimedia communication to this generalised approach to media (and service) delivery that we have sketched in this paper.

References

- [3GP04] 3GPP. Multimedia Broadcast/Multicast Service: Protocols and Codecs. 3GPP TS 26.346, 2004.
- [CD99] J. Clark, S. DeRose, XML Path Language (XPath) Version 1.0. W3C Recommendation, November 1999.
- [CRGMW96] Sudarshan S. Chawathe, A. Rajarama, H. Garcia-Molina, J. Widom, Change Detection in Hierarchically Structured Information. Proceedings of the ACM SIGMOD International Conference on Management of Data, June 1996.
- [(Ed04] B. Ramsdell (Ed.). Secure/Multipurpose Internet Mail Extensions (S/MIME) Version 3.1 Message Specification. RFC 3851, July 2004.
- [Gre05] J. Greifenberg, Generation, Encapsulation and Processing of Metadata with Internet Media Guides. Master's thesis, Universität Bremen, 2005.
- [Han96] M. Handley, The SDR Session Directory: An Mbone Conference Scheduling and Booking System. Technical report, University College London, April 1996.
- [HC04] H. Holbrook, B. Cain, Source-Specific Multicast for IP. Internet Draft draft-ietf-ssm-arch-06.txt, Work in progress, September 2004.
- [HCBO99] M. Handley, J. Crowcroft, C. Bormann, J. Ott, Very Large Conferences on the Internet: the Internet Multimedia Conferencing Architecture. Computer Networks, The International Journal of Computer and Telecommunications Networking, Special Issue on Internet Telephony, **31** (1999) 3, 191-204. Elsevier, North Holland, 1999.
- [HJ98] M. Handley, V. Jacobsen, Session Description Protocol. RFC 2327, April 1998.

- [Loo05] K. Loos, Papageno: Eine skalierbare Distribution-splattform für heterogene Empfängergruppen. Master's thesis, Universität Bremen, 2005.
- [LPP03] Juha-Pekka Luoma, J. Peltotalo, S. Peltotalo, MUPPET: Internet Media Guide Unidirectional Point-to-Multipoint Transport. Internet Draft draft-luoma-mmusic-img-muppet-04.txt, Work in progress, December 2003.
- [Moa97] R. Moats. URN Syntax. RFC 2141, May 1997.
- [NS04] Y. Nomura, H. Schulzrinne. SIP Event Notification for Internet Media Guides. Internet Draft draft-nomura-mmusic-img-notify-00.txt, Work in progress, July 2004.
- [NWL+04b] Y. Nomura, R. Walsh, Juha-Pekka Luoma, J. Ott, H. Schulzrinne. Requirements for Internet Media Guides. Internet Draft draft-ietf-mmusic-img-req-07.txt, Work in Progress, June 2004.
- [OL05] J. Ott, K. Loos: Using HTTP for IMG Transport, Internet Draft draft-ott-mmusic-img-http-00.txt, Work in Progress, October 2005.
- [The05] The TV-Anytime Forum. Specification Series: S-2 (Phase 2) On: System Description. Technical report, 2005. SP002v2.0.
- [WCPP05] H. Mehta, R. Walsh, I. D. D. Curcio, J. Peltotalo, S. Peltotalo. SDP Descriptors for FLUTE. Internet Draft draft-mehta-rmt-flute-sdp-05.txt, Work in Progress, January 2006.
- [WLP+05] R. Walsh, Juha-Pekka Luoma, J. Peltotalo, S. Peltotalo, J. Greifenberg. The IMG Envelope. Internet Draft draft-walsh-mmusic-img-envelope-03.txt, Work in progress, June 2005.



JÖRG OTT is Professor for Networking Technology with a focus on Protocols, Services and Telecommunications Software at the Networking Laboratory in the Electrical and Communications Engineering Department at Helsinki University of Technology. From 1997 through early 2005, he was Assistant Professor in the Computer Networks group at the Universität Bremen and member of the Center for Computing Technologies (TZI). He received his Doctor in Engineering (Dr.-Ing.) in 1997 from Technische Universität Berlin. His research interests are in protocol and system architectures for multipoint communications, content distribution, and multimedia conferencing, including IP Telephony as one special interest area. His current research focus is on communication in challenged networks, particularly on delay-tolerant networking. He has been active in the IETF since 1993 where he has been co-chair of the MMUSIC working group since 1997 and was co-chair of the SIP working group from its foundation in 1999 to October 2002.

DIRK KUTSCHER is a researcher at Technologiezentrum Informatik (TZI), Universität Bremen. He received his Doctor in Engineering (Dr.-Ing.) in 2003 from Universität Bremen. He has worked in several national and international research projects in the areas of wireless and interpersonal multimedia communications and has continuously participated in IETF standardization since 1998. He has co-authored two books, several Internet Drafts, RFCs and scientific publications. His current research focus is on architectures and communication protocols for challenged environments, and he is participating in the on-going research project Drive-thru Internet. In 2005, Dr. Kutscher has been working as a visiting researcher at KDDI R&D Labs. Inc. in Yokosuka Research Park, Japan, where he has investigated new approaches for simulating wireless communication protocols and media media-independent handover services for future heterogeneous networks.