

Defining a Standard for Orthodontic Electronic
Data
Considering HL7 for textual data

Antonio Magni
panio@antoniomagni.org

revision 1
January 16, 2006

Abstract

The American Dental Association Standard Committee for Dental Informatics is currently working on defining a standard for orthodontic electronic data. Health Level 7 (HL7) is a well established medical information standard which could be a candidate for representing such data. This paper aims to introduce HL7 to the interested parties and to expose the process of its refinement. The refined HL7 would allow the transfer of orthodontic electronic data using an already existing and well established standard. Binary data such as images are not part of HL7, which requires the cooperation with another standard: HL7 suggests the integration with DICOM to represent imaging data.

Contents

1	Introduction	3
1.1	The necessities of orthodontic data	3
1.2	Comparing HL7 to other standards	4
2	Health Level Seven (HL7)	4
2.1	The Organization	4
2.2	The Standard	5
2.2.1	Version 2.x (v2)	6
2.2.2	Version 3 (v3)	7
3	Building blocks of v3	8
3.1	Reference Information Model (RIM)	8
3.2	Domain Message Information Model (D-MIM)	9
3.3	R-MIM	10
3.4	Hierarchical Message Descriptors (HMD)	10
3.5	Message Types (MT)	10
3.6	Storyboards	10
3.7	Application Roles	11
3.8	Trigger Events	11
3.9	Interactions	12
4	Refining HL7	12
4.1	Project Initiation	13
4.2	Requirements Documentation	13
4.3	Specification Modeling	13
4.4	Specification Documentation	14
4.5	Specification Approval	14
4.6	Specification Publication	14
4.7	Implementation Profiling	15
5	HL7 and orthodontic electronic data	15
5.1	Implementing the standard in HL7	15
5.1.1	Project Initiation	16
5.1.2	Requirements Documentation	16
5.1.3	Specification Modelling	19
5.1.4	Specification Documentation	21
5.1.5	Specification Approval	21
5.1.6	Specification Publication	22
5.1.7	Implementation Profiling	22
5.2	Images and other binary data	22
6	Discussion	23
7	Conclusion	24

1 Introduction

In this paper we first present the context of the standard, then give an introduction to HL7 and its core elements, summarize what is needed to refine the standard and define a list of items that need to be taken care of while writing the digital orthodontic standard in HL7. It is not our intention to provide a solution, merely to provide the reader with a summary of what HL7 is and how it could be applied to orthodontic data.

Refining HL7 by adding an orthodontic domain alone cannot solve the problem of defining a standard for digital orthodontic data: orthodontic data includes images, which HL7 does not know how to handle. Nonetheless, HL7 integrates very well with other external standards, such as DICOM. In order to benefit from HL7, it would therefore be necessary to also refine an imaging standard to be able to accommodate orthodontic images. The integration of the two standards would ideally be coordinated by a higher level standard organization (such as the ADA SCDI or IHE) whose job would be to provide the community with documentation that specifies a standard way of using these lower level standards.

In the next section we discuss requirements that must be met in order to establish an orthodontic electronic data standard. This information is based on the meetings and discussions of the ADA SCDI working group 11.6.

1.1 The necessities of orthodontic data

Many software products have been developed for various orthodontic related tasks, ranging from clinical to administrative use. Yet none of them conform to a well accepted international standard.

Current digital dental standards do not adequately cover all the requirements of the orthodontic domain. Orthodontics has information elements related to treatment¹ and imaging that do not exist in the dental domain. A new standard that covers the structure formats and relationships of these additional elements is needed[?].

Orthodontists need common definitions, structures, formats and encodings of data relevant to the orthodontic treatment of patients, in a language- and platform-neutral and extensible fashion. There is also a demand for application-level protocols for the secure and reliable data exchange between applications and sites, with the goal of facilitating sharing and collaboration between patients, doctors, other therapy and service providers, etc[?].

Summarizing, the orthodontic electronic data standard should provide:

1. A way to easily exchange patient data between colleagues;
2. A way to easily port patient data from one system to another;

¹Diagnosis, treatment planning, outcomes analysis, appliances, root anatomy.

3. A way to share patient data between different software systems within the same practice, thus eliminating the need of having to maintain different sets of the same patient data;
4. A way to secure patient data, s.t. patient's rights will not be violated (i.e. implement HIPAA regulations);

To this extent, we shall analyse the HL7 standard to see if it can meet these requirements.

1.2 Comparing HL7 to other standards

This is a difficult task, as there are no standards as complete as HL7 to compare to. The only standard with similar goals is MEDIX, but we were not able to find any recent work on the Internet (most recent documentation is dated 1997). MEDIX stands for Medical Data Interchange and was being developed by IEEE. HL7 was closely collaborating with MEDIX in order to provide compatibility between the two.

While numerous standards for materials, equipment and techniques have been developed in the dental and orthodontic field, very few actually specify electronic data. Most work in this field has been done by the ADA. On the other hand, various electronic medical data standards have been developed over the past few decades: The American Standards Committee (ASC), American Society for Testing and Materials (ASTM) and Institute of Electrical and Electronic Engineers (IEEE) all have produced complete and usable standards. None of them, though offer an integrated solution: for example, the ASC X12 is a standard for business documents only, while the ASTM standards defines limited domains, such as electronic health records, authentication of health care information, universal healthcare identifier properties, users authentication amongst others.

HL7 took some of these standards, worked together with their committees and formed a new, integral standard that offered a complete solution for the medical environment. Now in its third major version, it offers the most elaborate and modern standard currently available for the medical field.

2 Health Level Seven (HL7)

HL7 is the acronym for both an organization and the standard that the organization supports and maintains. In this section we provide an introduction to HL7 as an organization and a standard.

2.1 The Organization

Founded in 1987, Health Level Seven is an American National Standards Institute (ANSI)-Accredited Standards Developing Organization (SDO) which focuses on the electronic interchange of clinical, financial and administrative in-

formation among independent healthcare-oriented computer systems. HL7 is a not-for-profit volunteer organization whose members are providers, vendors, consultants, government groups and others who have an interest in the development of healthcare standards. According to HL7[?, p. 2], 90% of healthcare system vendors are members, comprising over 2,200 health industry members. HL7 was designated by the ANSI as an ANSI-accredited SDO and since then has published and received ANSI approval for various medical standards. HL7 collaborates with 14 other standard committees and has affiliates in 27 different countries.

The organization is well structured and subdivided in 26 Technical Committees (TC) and 18 Special Interest Groups (SIG). TCs focus on the creation, maintenance and extension of the HL7 Protocol Specifications, each of which specializing on a different subject matter. SIGs are concerned with projects that aid the application and implementation of the standard itself. The Java SIG, for example, is concerned with developing a Java Application Programming Interface (API) to the HL7 information model.

The goal of HL7 is to specify a standard way for programs of different vendors to be able to communicate by easily exchanging patient data. It focuses on the communication level, defining how to build a well formed message that can be read by HL7 compliant systems. It does not cover the format in which data should be stored. The documentation is very specific and makes frequent use of pseudo-code, flow charts and UML diagrams (for example to aid the reader in understanding how to compose a well formed HL7 message).

Currently, the members of HL7 are working on version 3 of their standard. They claim it to be substantially more advanced and more complete than their previous version 2. The latest ANSI approved version, however, is still 2.5.

2.2 The Standard

As a standard, HL7 is an application protocol for electronic data exchange in healthcare environments. The number seven refers to the highest level of the International Organization for Standardization (ISO) communications model for Open Systems Interconnection (OSI) - the application level. The application level addresses definition of the data to be exchanged, the timing of the interchange, and the communication of certain errors to the application. The seventh level supports such functions as security checks, participant identification, availability checks, exchange mechanism negotiations and, most importantly, data exchange structuring.

The HL7 developed documentation to structure medical information in a universal and coherent fashion. The different object classifications and subdivisions are mainly oriented towards textual information. That is, HL7 did not deem necessary to define how to encode or store binary data (i.e. digital images, movies, audio files, 3-dimensional volumes, ...). Many different encoding formats for such data have already been developed, making it unnecessary for HL7 to define a new one. Instead, HL7 created a well organized structure for all medical knowledge, including a placeholder for binary data. Thus, if it were

necessary to send a patient's X-ray over the network using HL7 messaging, it would first be necessary to choose which format to save and send our X-ray in (hopefully picking one that is readable at the receiver's end!). Next the X-ray would be wrapped by a newly created HL7 message. At this point the X-ray could be transferred between HL7 compliant systems.

Although HL7's latest approved ANSI standard is version 2.5, over the past 14 years HL7 has been working on version 3, which is still in a ballot state.

2.2.1 Version 2.x (v2)

ANSI has approved the first version of HL7, v2.1 back in 1991: currently version 2.5 is the latest ANSI approved² version. It specifies a set of rules on how to organize medical data so that it can be sent and received across networks and removable devices efficiently and reliably. This is all taken care of at the application level only, as basic network operation, such as error control, character conversion and message length are assumed to already be taken care.

Data is transferred using atomic units called *messages*. Within each message, data *fields* are grouped in *segments*. Each field is nothing but a string of characters with attributes³ associated with it.

Version 2.x is a simple static documentation of messages and their fields. It does not define interactions between processes and actors. In addition is not encoded using XML. It uses a system called *vertical bars*⁴ instead. However, HL7 has released a document on how to encode HL7 messaging version 2.x using XML[?].

Although widely⁵ used and still an international standard, Version 2.x is being phased out. HL7 thought it was necessary to make major improvements that would require a complete reorganization of the methodology used to develop the specification. To better understand why the organization began a new, completely revised edition of their standard, I shall quote a part of the Introduction of the manual[?]:

The HL7 v2.x development process is entirely ad-hoc. There is no explicit methodology. Members receive no formal guidance in constructing messages. Trigger events and data fields are described solely in natural language. The structural relationships among data fields are not clear. Segments are reused in many messages and message definitions are reused for many trigger events. In order to accommodate this extensive reuse, most data fields are optional. Chapters are inconsistent in their use of trigger events versus status

²Approved on June 23, 2003.

³These are position, max length, data type, optionality, repetition, table, ID number and name.

⁴Vertical bars (a.k.a. "pipe" character) are used to separate the fields in the messages. Hence the name.

⁵According to HL7 board members, 90% of hospitals in the USA make use of some kind of implementation of v2.x

codes. There is no specification as to when a specific kind of health-care information system should be expected to honor a trigger event or accept a message.

With v2.x, a Technical Committee creates messages by editing word processing documents directly. The meta-data is not available in a structured form until the staff and volunteers tediously extract it from the word processing documents after publication.

In summary, there is substantial need to improve this old process in order to handle the breadth and complexity of the challenges HL7 faces today. Our industry will benefit because this new process results in a more rigorous specification.

In addition, Version 2.x does not take care of security and patient confidentiality[?, p. 1-13] and is silent on messages to support the integration of a patient's health record across multiple delivery entities of a healthcare delivery system. This would also include messages to insure central control and integrity of information that was "merged" between multiple delivery entities. [?, p. 1-15]

Despite the above mentioned drawbacks, development of v2.x still continues. HL7 is working on 2.6 and some Special Interests Groups of HL7 are talking about 2.7. This is because it has been widely implemented and it works. The transition to HL7 v3 would be very costly, and to the front-end user, it would not show any apparent change.

2.2.2 Version 3 (v3)

In 1992 HL7 made a fundamental shift in the methodology used to develop its standard specifications. The new methodology, referred to as HL7 Version 3.0 or just v3, is a model-driven methodology based upon modern object-oriented software and domain modelling practices. The v3 project represents a new approach to clinical information exchange. It is built from the ground up around a single object model, the Reference Information Model (RIM) and a rigorous UML based methodology that ties model to messages and finally to the message's expression in XML syntax.

The v3 specification is built around subject domains, for each of which it provides storyboards descriptions, trigger events, interaction designs, domain object models (derived from the RIM), hierarchical message descriptors (HMDs) and prose description of each element. Implementation of these domains further depends upon a non-normative v3 Guide and normative specifications for data types, the XML Implementable Technical Specifications (ITS) (or message wire format), message control wrappers and transport protocol. HL7 v3 is the most definitive HL7 standard thus far, incorporating more trigger events and message formats than any previous version.

HL7's primary goal for v3 is to offer a standard that is definite and testable, and to provide certification of vendor's and implementer's conformance. Thus the development principles behind v3 lead to a more robust, fully specified standard.

Effectively, the main difference between v2.x and v3 is that v3 specifies actors and processes and the relationship between them along with how to send and encode messages. v2.x only deals the messages themselves.

As of the time of this writing⁶, v3 is still in a ballot state, although it has just released its first revision (the documentation was shipped out to members in November 2005).

Further discussion and analysis will be limited to v3 only. In our opinion the orthodontic community would not benefit from using v2.x.

3 Building blocks of v3

In order to make any use of the standard, it is essential to first understand its basics. v3 is not a simple read: its concepts are abstract and complex. In this section we present the reader with the elements of v3 needed to understand the refinement process described in Section 4. It is not our intention to provide a complete description of HL7 v3.

HL7 messaging components can be of two types: static and dynamic. Dynamic components describe interactions between systems (Storyboards (Sec 3.6), application roles (Sec 3.7), trigger events (Sec 3.8) and interactions), while static components describe the static content of the messages (D-MIM (Sec 3.2), R-MIM (Sec 3.3), HMD (Sec) and MT (Sec)).

Static components

3.1 Reference Information Model (RIM)

At the heart of v3 there are *Information Model Components* such as the *Reference Information Model*. The RIM serves as a common source of information for the entire specification. It provides an explicit representation of the semantic and lexical connections that exist between the information carried in the fields of HL7 messages. The RIM is used to express the information content for the collective work of the HL7 Working Group. It is the information model that encompasses the HL7 domain of interest as a whole[?, Sec. 2.2.2]. In other words the RIM provides a means of specifying the information content of messages through a common information model that clarifies the definitions and ensures that they are used consistently across all v3 messages defined by all Technical Committees.

The HL7 RIM is a critical component of the v3 development process. It is the root of all information models and structures developed as part of the v3 development process. It is the heart of the HL7 v3 standard.

The RIM is documented using pure UML diagrams. The RIM is built in an

⁶December 2005.

object-oriented fashion, making use of *classes*⁷, *generalizations*⁸, *associations*⁹, *data types*¹⁰ and *attributes*. Class attributes are the core components of the information model. They are the source for all the information content of HL7. The RIM contains six core classes (Act (Actions), Entity (People, Places and Things), Role, Act_Relationship (connects Acts), Participation (connects Roles to Acts), Role_Link (connects Roles)) from which all other classes derive.

The majority of attributes are descriptive in nature. All of these elements are controlled by *constraints*¹¹ and *vocabulary*¹².

Furthermore, the documentation contains a whole section that defines the implementation technology, called ITS or *Implementation Technology Specification*. This part defines how to represent RIM objects for transmission over some kind of media (email, CD, removable disks, ...). This is the most low-level definition of the standard as it descends into ISO levels 6 and 5. HL7 has adopted XML for its initially balloted ITS, and has selected the XML schema recommendation as the best method within the XML family of standards.

3.2 Domain Message Information Model (D-MIM)

Like the other models included in the v3 documents, the Domain Message Information Model (D-MIM) is a diagram that shows the relationships between the classes. Differently from the RIM, though, the D-MIM (as well as the R-MIM) make use of a modified UML. The D-MIM is a subset of the RIM (see Section 3.1) that includes a fully expanded set of classes (always clones of RIM classes), attributes and relationships that are used to create messages for any particular domain. For example, the set of classes that are used by the Medical Records/Structured Documents domain is quite different from that used by the Patient Administration domain. The D-MIMs for these two domains, then, will be quite different, although both will be derived from the RIM.

The D-MIM provides a solution to the information requirements of a particular problem domain. The mapping of the requirement's domain information model to the RIM is used to identify which RIM classes need to be included in the D-MIM. In some cases it may be necessary to include multiple clones of the same RIM class. Each clone is given a unique name that is reflective of its business use.

⁷A *class* is an abstraction of things or concepts that are subject of interest in a given application domain. Classes are the people, places, roles, things, and events about which information is kept. Classes have a name, description, and sets of attributes. Instances of classes are called *objects*.

⁸A generalization relationship is a connection between classes (as opposed to objects).

⁹An association defines a relationship between objects.

¹⁰Data types are the basic building block of attributes. They define the structural format of the data carried in the attribute and influence the set of allowable values an attribute may assume.

¹¹Constraints narrow the set of possible values that an attribute can take on.

¹²A vocabulary domain specifies all valid values in an instance of a field or attribute.

3.3 R-MIM

Refined Message Information Models (R-MIMs) are used to express the information content for one or more HMDs that originate from the root class identified by the Entry Point in the R-MIM. Each R-MIM is a subset of the D-MIM and contains only those classes, attributes and associations required to compose the set of messages derived from the HMDs that originate from the R-MIM root class. Classes, attributes and associations that are not required for those HMDs are eliminated and the generalization hierarchies are also collapsed.

3.4 Hierarchical Message Descriptors (HMD)

In simplest terms, an HMD is a tabular representation of the sequence of elements (i.e., classes, attributes and associations) that define the message without reference to the implementation technology. The HMD defines a single base message structure - the "common" message type. This base message structure is never sent and thus has no corresponding trigger event. It is the template from which the other specific and corresponding message types are drawn. The HMD and its contained message types may be represented as a spreadsheet.

3.5 Message Types (MT)

A message type represents a unique set of constraints applied against the common message.

Dynamic Components

3.6 Storyboards

While reading through the HL7 specifications and other resources, we encountered various ways of defining storyboards, each of them pointing to the same or similar definition. Since the term "storyboard" is not common within standard development, we thought the reader could benefit from the following list:

- A Storyboard details a temporally sequenced series of actions/interactions involving one to many participating entities (e.g. human and/or system), and may, over its course, provide specific value to one or more of the involved entities.
- A Storyboard is a plain language description of a series of steps involving some exchange of information between different participants to achieve the objectives of a healthcare business process. The list of steps can be in generalized, abstract terms, or in the form of a real-world example.
- The Storyboard answers the question "for what purpose is this information being shared?"
- Storyboards are a means of providing context to the definitions of trigger events(Section 3.8).

- The process of storyboarding lays the foundation for describing HL7 messages and their content.
- A storyboard narrative is a description of a real-life event that provides the necessary context for the development of a specific interaction described in the storyboard.
- A storyboard consists of a short description of its purpose and an interaction diagram that shows the progression of interactions between the application roles (see Section 3.7).
- Storyboards are informative as opposed to normative: they exist to clarify other normative sections of the standard.

The storyboard concept is borrowed from the movie and animation industry, and is useful to the development of HL7 messages for the same reasons proven in that industry:

- A storyboard depicts a story using a series of “snapshots” or events in chronological sequence;
- Each snapshot represents a recognizable, meaningful moment in the sequence of events that the reader must know about to understand the overall sequence and result;
- Each snapshot illustrates the key participants in the storyboard and their interaction with other players;
- The whole series of snapshots provides a coherent description of a complete process or activity.

3.7 Application Roles

Application roles represent a set of communication responsibilities that might be implemented by an application. Thus they describe system components or sub-components that send and/or receive interactions. Practically, an application role represents a computer or a program that plays a role within the scenario of sending/receiving HL7 messages.

3.8 Trigger Events

A trigger event is an explicit set of conditions that initiate the transfer of information between system components (application roles). It is a real-world event such as the placing of a laboratory order or drug order. In the v3 standard, trig-

ger events are one of interaction¹³, state-transition¹⁴ or user-request trigger¹⁵. Most trigger events are State-Transition based and will be encountered when reading the dynamic message information model (D-MIM) defined to support a particular message interaction.

3.9 Interactions

Interactions are used to explicitly define interactions between application roles. It is a unique, one-way transfer of information.

A single Interaction explicitly answers the questions: 1. Which type of system component sends a particular type of message; 2. To what type of receiving system component the message type is sent; 3. How a system knows when to send a particular type of message; 4. What the particular message type is;

Interactions are typically represented using interactions diagrams (Fig. 5) and used within storyboards.

4 Refining HL7

In this section we discuss the steps involved in adding new domains to the HL7 specifications.

Currently, HL7 does not provide a dental or orthodontic domain. Therefore, if we would like to implement the orthodontic electronic data standard using HL7 we must learn how to refine and expand the current HL7 specifications.

The HL7 Development Framework(HDF)[?] is a document which details the processes of the HL7 development methodology. The full process of creating an HL7 specification is divided in seven steps[?]:

1. Project Initiation
2. Requirements Documentation
3. Specification Modelling
4. Specification Documentation
5. Specification Approval
6. Specification Publication
7. Implementation Profiling

¹³Trigger events can be based on another interaction. For example, the response to a query (which is an interaction) is an Interaction Based trigger event.

¹⁴Trigger events resulting from a state transition as depicted in the State Transition Model for a particular message interaction. The trigger for canceling a document, for example, may be considered a State Transition Based trigger event

¹⁵Trigger events may be based on a user request. For example, the trigger event that prompts a system to send all accumulated data to a tracking system every 12 hours is considered User Based.

The following subsections were copied as is from the HDF. The HDF contains individual chapter for each of the seven steps listed above. Please refer to the HDF[?] for further details, examples, tools and templates.

4.1 Project Initiation

During project initiation the project is defined, a project plan is produced, and project approval is obtained. The primary deliverable produced during project initiation is the project charter. The objectives of the project charter are to:

1. Define project scope, objectives, and intended deliverables
2. Identify project stakeholders, participants, and required resources
3. Document project assumptions, constraints, and risk
4. Prepare preliminary project plan and document inter-project dependencies
5. Obtain project approval and launch the project

4.2 Requirements Documentation

During requirements documentation the problem domain is defined, a model of the domain is produced, and the problem domain model is harmonized with HL7 reference models. The primary deliverable produced during requirements documentation is the requirements specification. The sequence of steps to create the requirements specification are:

1. Document Business Process: Dynamic Behavior and Static Structure
2. Capture Process Flow: UML Activity Diagram
3. Capture Structure: Domain Analysis Model and Glossary
4. Capture Business Rules: Relationships, Triggers, and Constraints
5. Harmonize the Domain Analysis Model with HL7 Reference Models

4.3 Specification Modeling

During specification modeling reference models are constrained into design models through a process of iterative refinement driven by requirements specifications and following specification design rules, conventions, and guidelines. The primary deliverable produced during specification modeling is a set of specification design models (D-MIMs). The steps are:

1. Build design models of static information views
2. Construct design models of behavioral views

3. Define reusable design model components
4. Construct design models of collaboration and interaction
5. Harmonize design models with HL7 Reference Models

4.4 Specification Documentation

During specification documentation the specification design models are packaged into logical units, supplemented with explanatory text, and prepared for approval. The primary deliverable produced during specification documentation is a proposed specification. The steps to produce specification documentation are:

1. Organize design model elements into logical packages
2. Compose explanatory text, examples, and design rationale
3. Update design models and requirement specifications
4. Assemble a proposed specification package
5. Submit specification for approval

4.5 Specification Approval

During specification approval the proposed specification is subjected to a series of approval steps. The specific approval steps vary by kind of specification, level of approval, and realm of interest. The primary deliverable produced during specification approval is an approved specification. The approval steps are:

1. Obtain TSC and Board approval to ballot specification
2. Form a ballot pool and conduct specification ballot
3. Assess negative ballots and affirmative comments
4. Modify specification in response to ballot comments
5. Resolve negative ballot responses and if necessary re-ballot

4.6 Specification Publication

During specification publication the approved specification is prepared for prepared for publication and distribution. The primary deliverable produced during specification publication is a published specification. The steps to publication are:

1. Obtain TSC and Board approval to publish specification
2. Prepare specification for publication

3. Submit publication to standards authorities (ANSI/ISO)
4. Render the specification in various forms of publication media
5. Post and distribute approved specifications

4.7 Implementation Profiling

During specification profiling specification, models are further refined and specifications furthered constrained. This refinement and constraining follows the same set of design rules, conventions, and guidelines used in the development of the specification to produce a profile of the specification for use in a particular environment by a defined community of users. The primary deliverable produced during specification profiling is a set of specification profiles and conformance statements. The steps to produce these profiles and conformance statements include:

1. Identify community of uses for published specification
2. Further refine and constrain specification design models
3. Document exceptions, extensions, and annotations to specifications
4. Prepare and publish specification profile
5. Prepare and publish conformance statements

5 HL7 and orthodontic electronic data

This section is devoted to the evaluation of refining HL7 with respect to orthodontic electronic data.

After analysing HL7, we realized that it provides a stable framework to aid the development of an orthodontics electronic health record standard. But what would it mean, specifically, to implement our standard using HL7? How much work would ADA SCDI WG 11.6 need to accomplish?

5.1 Implementing the standard in HL7

This section contains a summary of the process of refining HL7 for orthodontics. It is intended to give the reader an idea of what it would be like to follow this path. Please refer to the HL7 Development Framework[?] for more details on the process.

HL7 defines the exchange of messages between applications. Using HL7 to define an orthodontic data standard would mean looking at orthodontic data from a communication perspective. As an example HL7 looks at a patients first visit as information that needs to move from one system (application role) to another. The patient data would be exported and imported from and to applications within messages.

The Modelling and Methodology technical committee at HL7 has developed a manual[?] which details the processes of the HL7 development methodology. In general, refinement of HL7 starts from the creation of storyboards. Orthodontic experts should write a set of storyboards, one for each exchange of clinical and financial information. These can then be used to create an orthodontic D-MIM and R-MIMs. From R-MIMs HMDs and message types can be created.

According to the HDF, the full process of creating an HL7 specification is divided in seven steps (see Section 4). All seven steps require, within the working group, the presence of orthodontics domain experts, HL7 facilitators and memberships, and an HL7 expert.

In the next sections we review the seven steps listed in Section 4 applying them to the particular task of developing an orthodontic electronic standard.

5.1.1 Project Initiation

This entails getting a consensus from the working group to go ahead and use HL7 for orthodontics. A project charter has already been produced by the ADA SCDI WG11.6 ex co-chair Philip DeSmedt at Align Technology and should be available off of the ADA SCDI web site. It needs to be adapted and focused towards creating a new HL7 domain, according to HL7 regulations.

5.1.2 Requirements Documentation

These are the steps involved in delivering a **requirements specification**.

Storyboards The first step in the requirements gathering process is to develop a description of the orthodontic electronic data exchange problem. This is done using Storyboards (see Section 3.6).

Documentation of the orthodontic electronic data processes involves unambiguously describing both the structure and the behavior/function of the entities involved in the processes. This documentation should be created making use of the knowledge of orthodontic professionals and the project charter, and should be captured in a Storyboard.

Activity Diagram Currently¹⁶ ADA SCDI working group 11.6 is partially working on this by defining an information model framework of the orthodontic data domain. This work, though, would need to be steered toward the creation of Storyboards. With storyboards, it will be possible to expand them into activity diagrams¹⁷ (see Figure 1).

¹⁶As of October 5th, 2005.

¹⁷Activity Diagram is defined in version 1.4 of the UML - "An Activity (Graph) Diagram is a variation of a state machine in which the states represent the performance of actions or sub-activities, and the transitions are triggered by the completion of the actions or sub-activities. It (therefore) represents the state machine of a procedure (or process) itself. The purpose of this diagram is to focus on the flows driven by internal processing (within a system or subsystem)."




Figure 1: An Activity Diagram identifies a sequence of steps and the information that is transferred from one participating role to another. Sometimes called a “Swim-lane Diagram”, the pictures represent the flow of control among the steps and help identify what information is required to be transmitted to achieve the objectives of the Storyboard. Of particular interest, is the data information exchange focus of HL7, are the Activity Diagram semantics that depict the passing of objects (e.g., data, information, messages, documents) between swimlanes.

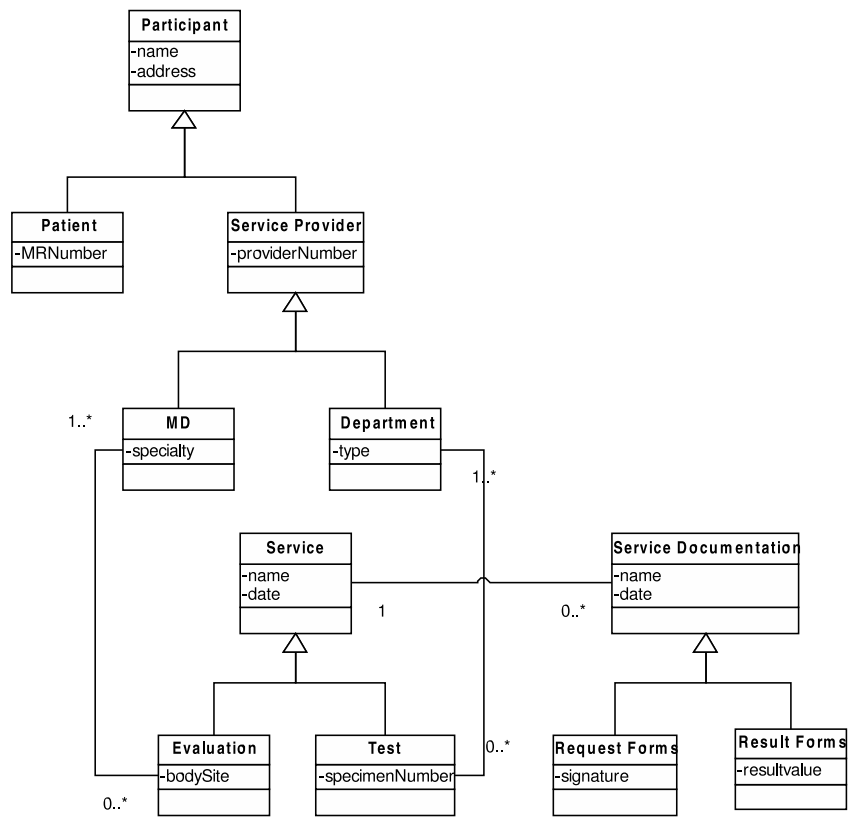


Figure 2: The Domain Analysis Model describes the key information needed to be shared to achieve the objectives of the Storyboard.

Domain Analysis Model and Glossary It will then be necessary to develop a Domain Analysis Model using a UML Class Diagram. The class diagram simply needs to identify the domain concepts-of-interest and their static inter-relationships using UML's tools¹⁸. It does not need to be fully implemented (i.e. ready to be translated into code), with all methods and attributes (see Figure 2).

The glossary is needed to clarify the terms used by orthodontic professionals to identify the processes themselves. A draft of orthodontic data types could be a useful resource for this task and can be found in [?]. The glossary should ultimately be in the form of a two-column table: Term vs. Definition.

Relationships, Triggers and Constraints With the activity diagram, domain analysis model and domain glossary it will be possible to carefully describe the structure of the data/information to be exchanged. This is then added to the activity diagram using the object/instance iconography¹⁹.

Harmonization Finally it will be required to harmonize the artifacts developed in the preceding steps with the existing HL7 reference models. To do so, any inconsistencies, redundancies and omissions must be aligned.

5.1.3 Specification Modelling

For the specification modelling, the deliverable is a specification design model, i.e. a set of D-MIMs (refer to Section 3.2). This is another five step process, similar to the one needed to deliver a requirements specification. At this stage it is necessary to be more detailed, in order to be able to produce an HL7 compliant specification model. This entails looking for an already existing D-MIM in the HL7 standard that somewhat meets the requirements specifications. If one already exists, it should be modified by adjusting the class clone names, attributes and relationships. Otherwise a new one should be created by cloning already existing classes from the RIM.

In order to do so, it will be necessary to:

- iteratively refine the activity diagrams from the requirements specifications using UML Sequence, Collaboration and State Transition Diagrams;
- construct collaboration diagrams from the system responsibilities for sending and receiving information(see Fig. 3);
- construct a sequence diagram to show the set of interactions between the application roles in the sequence required to meet the objectives of the Storyboard (see Fig. 5);
- translate the glossary into a Vocabulary Specification Schematic.

¹⁸Associations, association names and multiplicities.

¹⁹The object/instance iconography is an object-oriented programming terminology where an object is an abstract thing (i.e. a Pine Tree) and an instance is one specific object (i.e. The pine tree located at the corner of 12th St and Houston in Cleveland, OH).

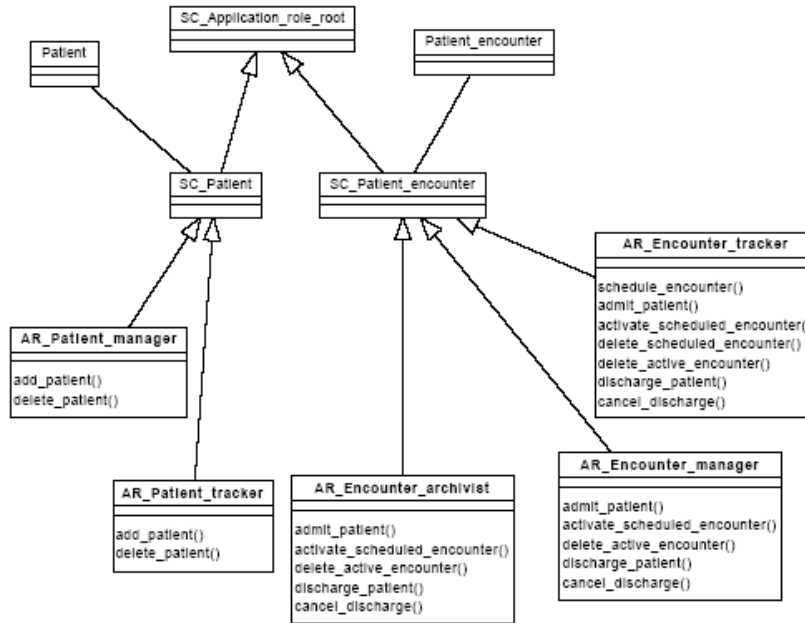


Figure 3: An example of a UML Collaboration Diagram.

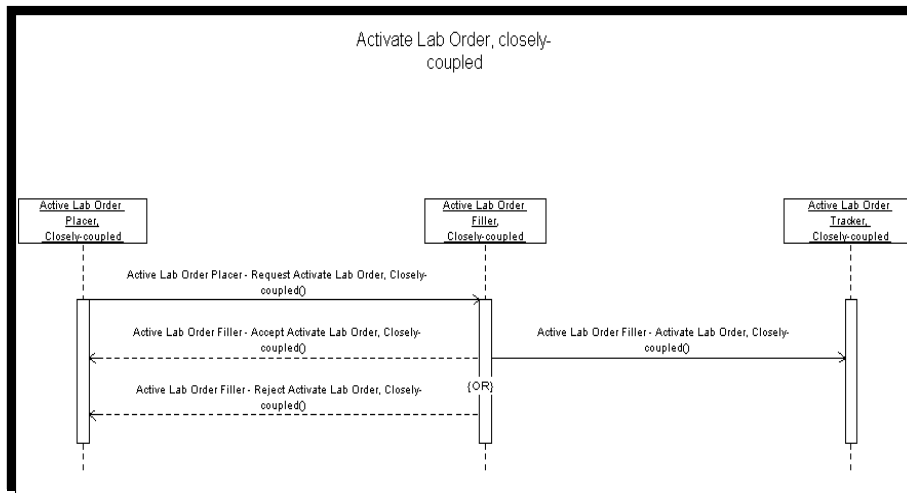


Figure 4: The UML Sequence Diagram details an interaction, i.e. specific trigger event, sending application role, receiving application role, receiver responsibility and optionally the interactions the receiving application must initiate.




Figure 5: An example of a Vocabulary Specification Schematic.

Finally, it will be necessary to harmonize the design models with HL7 reference models. It is advisable to make extensive use of the suggested tools listed in the HDF[?] for this process.

5.1.4 Specification Documentation

For the specification documentation, the primary deliverable is a proposed specification. This step focuses on documenting what has been designed so far. This includes:

- Using correct naming according to artifact naming convention defined in the HL7 v3;
- Writing explanatory text, design rationale and examples for each design artifact;
- Making sure there aren't any inconsistencies in the design models (Section 5.1.3) and requirements specification (Section 5.1.2) ;
- Create and test links to all referenced files;
- Zip and submit the package for approval;

5.1.5 Specification Approval

During specification approval, the proposed specification is subjected to a series of approval steps, which may vary by kind of specification, level of approval, and realm of interest. The general steps include obtaining approval to ballot specification from technical committee and board, conducting a ballot, assessing ballots, modifying specifications and resolving negative ballots by eventually re-balloting.

The primary deliverable is an approved specification.

5.1.6 Specification Publication

Finally the specification is ready for publication. This last step produces as a primary deliverable a published orthodontics HL7 specification. This entails obtaining approval to publish from board and technical subcommittee, preparing and submitting publication to standard authorities (ANSI/ISO), post and distribute the specification in various forms of publication media.

5.1.7 Implementation Profiling

Once the specification is finished, it is ready to be implemented and/or further refined. This step involves performing tasks to directly aid the implementation of the specification by further constraining it for a variety of purposes, depending on the user. For example the orthodontic specification could be further constrained to conform with local medical laws, or local practice procedures. In addition, this step produces conformance statements used by system sponsors such as vendors to communicate to a user how their products meet HL7 specifications.

It is also possible to define sub-domains through constraints. This means that the orthodontic HL7 specifications could refine an already existing sub-domain, instead of the RIM directly.

Implementation profiling must also undergo HL7 approval.

5.2 Images and other binary data

Images are one of the most used elements of an orthodontic electronic health record. Although HL7 only specifies how to send/receive *textual* data, it does leave space for any kind of binary data. Implementers can therefore choose, through MIME²⁰ types, how to encode their data with flexibility. HL7 suggests the use of DICOM, in which case the patient's image would be stored in the patient record in DICOM format using the Encapsulate Data (ED) data type[?, in Foundation, Data Types, Sec. 2.4]. The HL7 organization has a special interest group (SIG) that specifically focuses on this issue. The Imaging Integration SIG works closely with DICOM (their members are also members of DICOM Working Group 20): their scope is to make sure that DICOM objects can be transferred through HL7 messages without data loss or errors.

Hence to implement images we would need to:

1. Define the orthodontic image domain (in storyboards, D-MIMs and R-MIMs);
2. Define how to encode image data;
 - (a) In the case of DICOM, the DICOM standard would need to be revised to be able to accommodate for orthodontic data.

²⁰A set of rules that defines how to send binary data (audio, images, movies,...) through email. This same set of rules can be used for any kind of information transmission, not just email.

3. Pick the HL7 message that best represents the use case and define how to store orthodontic images as an ED data type;

This would entail using an image format that is capable of optimally storing orthodontic images. Choosing HL7 as the basis for the Orthodontic Standard would therefore not alleviate us from the task of defining the image encoding.

6 Discussion

HL7 is a well developed, widely used medical specification framework that focuses on medical data transmission over electronic medium. It already includes a wide variety of medical domains, and is working towards including more. Although it is not sufficient to define the entire orthodontic domain, it allows for the integration of any number of other standards, so to accommodate for the needs of even the most complex domains.

After evaluating HL7, we summarized pros and cons into the following lists.

Pros

1. HL7, being a communications standard, provides a way to easily exchange patient data between colleagues, provided the colleagues make use of HL7 certified software systems;
2. HL7 provides a way to easily port patient data from one system to another, provided both systems implement HL7;
3. HL7 provides a way to share patient data between different software systems within the same practice, provided the different softwares implement HL7;
4. HL7 v3 provides specifications for conformance testing.
5. HL7 provides a distribution schema to popularize their specifications.
6. HL7 is a well established and spread out standard (v3 already implemented, even if still in ballot state!).
7. HL7 is complete, as it already specifies most medical processes (patient records, medical documents, financial and insurance documents...).
8. HL7 provides detailed documentation for extending, refining or adding domains to the standard.
9. HL7 is fully compatible with DICOM.
10. HL7 has a large community of developers and users.
11. HL7 has an active community, with responsive mailing lists.
12. HL7 is the only accredited standard to provide all these features.

Cons

HL7 does not specify a way to secure patient data This task is left for the software vendors to implement, but does not constitute a major problem: each message needs only to be wrapped in an encrypted transaction such as the HIPAA 275 transaction.

HL7 does not specify a way to encode images It only specifies how to send them. Non-textual data must be stored in an external format (e.g., DICOM) before being encapsulated into HL7 messages. On the other hand, being fully compatible with any kind of image encoding scheme, it provides an extra level of flexibility. From the ADA SCDI WG 11.6 standpoint, this means research and refining yet another standard.

HL7 is very complex in nature Mastering HL7 requires a good knowledge of UML, object-oriented concepts and modelling tools. All the different diagrams, definitions, classes, models and domains can cause the HL7 learning curve to be steep. Nonetheless, once mastered, its strict and organized nature makes it straightforward to manage and modify.

As mentioned above, images are beyond HL7's scope. They could be stored in their raw form within HL7 messages using the Encapsulated Data (ED) fields. Nonetheless, this would require defining a whole new set of image related attributes withing HL7 messages. We therefore suggest making use of a separate imaging standard. DICOM is a good candidate as its compatibility with HL7 is constantly being monitored by HL7 SIG *Imaging Integration* and DICOM WG 20 *Integration of Imaging and Information Systems*. Images could then be used in conjunction with HL7 messages to form the complete orthodontic domain.

7 Conclusion

According to our evaluation, HL7 is an adequate framework for defining the specifications for orthodontic electronic data. We advise making direct use of HL7 by refining it to accommodate for the orthodontic necessities. Nonetheless, if the ADA SCDI working group 11.6 considers this not an ideal approach, the HL7 documentation remains a rich and useful resource for the development of our standard.

Immediate future work should be geared towards developing HL7 storyboard-like documentation. Domain specialists could use already existing HL7 storyboards as examples to deliver a complete set of situations in which the transfer of clinical and financial information would be necessary. This documentation will greatly aid the definition of the orthodontic electronic data domain. The necessary modelling would then follow.

A Glossary

ADA American Dental Association.

ASC American Standards Committee.

ASTM American Society for Testing and Materials.

Artifacts Within the HL7 v3 standards the components that make up the documentation are each referred to as *artifacts*. This includes, storyboards, application roles, trigger events, D-MIMs, R-MIMs, HMDs, message types and interactions.

CMET Common Message Element Type.

DIM Same as D-MIM, probably an older acronym.

D-MIM Domain Message Information Model.

DSTU

HDF HL7 Development Framework.

HL7 Health Level Seven.

HIPAA Health Insurance Portability and Accountability Act

HMD Hierarchical Message Descriptions: A common description of the exact fields of a message and their grouping, sequence, optionality, and cardinality.

IEEE Institute of Electrical and Electronics Engineers.

ITS Implementation Technology Specifications: Separate syntax specifications, describing the algorithms used to encode and transmit the messages in an XML based character stream syntax.

MIME Multipurpose Internet Mail Extensions.

R-MIM Refined Message Information Model.

SCDI Standards Committee for Dental Informatics.

References

- [1] William E Harrell, Sharon Stanford, and Paul Bralower. ADA initiates development of orthodontic informatic standards. *American Journal of Orthodontics and Dentofacial Orthopedics*, August:153–156, 2005.
- [2] Inc. Health Level Seven. HL7 messaging standard version 2.5. an application protocol for electrical data exchange for healthcare environments. 2003.
- [3] Align Technology Inc., 3M Unitek, OraMetrix, and PracticeWorks. Plan for digital orthodontic standards. 2004.
- [4] The Health Level Seven. HL7 version 2: Xml encoding syntax, release 1. *ANSI*, 2003.
- [5] The Health Level Seven. *The HL7 Development Framework*. The Health Level 7, 2.0.0 edition, May 2005.
- [6] The Health Level Seven. HL7 v3 guide. 2005.
- [7] The Health Level Seven. Overview of health level seven. Mailed to Members, 05 2005.