Medical Infobahn for Europe
A. Hasman et al. (Eds.)
IOS Press, 2000

DICOM Extensions for Narrow-Band Networks

Jörg Riesmeier[a], Marco Eichelberg[a], Didier Lemoine[b],
Vytenis Punys[c], Nándor Balogh[d], Peter Jensch[e]
[a] *OFFIS e.V., 26121 Oldenburg, Germany*
[b] *ETIAM S.A., 35000 Rennes, France*
[c] *Kaunas University of Technology, 3006 Kaunas, Lithuania*
[d] *CorPuSNet Alapitvany, 1529 Budapest, Hungary*
[e] *Carl von Ossietzky Universität, 26111 Oldenburg, Germany*

Abstract. DICOM is today's de-facto standard for exchanging medical images. Since new image acquisition devices produce more and more image and non-image data, image compression has become an important part of the standard. However, the compression of non-pixel data also stored in DICOM data sets has been disregarded up to now. In the scope of an EU research project we have examined a large amount of real-world DICOM images to test whether or not there is a potential for compressing the non-pixel attributes. Especially for use with narrow-band networks extensions as proposed in this paper could be a solution to save valuable bandwidth.

1. Introduction

DICOM (Digital Imaging and Communications in Medicine) [1] is today's de-facto standard for exchanging medical images between devices from different vendors. In the last couple of years the development of DICOM not only followed the new medical image acquisition techniques producing growing amounts of pixel data but also tried to cover the emerging workflow management services.

The DICOM standard already supports various compression algorithms to reduce the size of large images, but compression was not defined under special consideration of narrow-band networks. Therefore, it makes sense to investigate whether there are any potential enhancements that could improve this situation. In particular, objects with large amounts of non-pixel data promise good results, because this textual information remains completely uncompressed in the current version of the standard. New DICOM objects containing only non-pixel data are, for example, Structured Reporting [2], Softcopy Presentation States [3] and Waveforms [4].

The results presented in this paper have been acquired in the scope of an European Union (EU) research project called SAMTA (Open Scaleable Architecture for Multimedia Telemedicine Applications) [5]. The goal of this project is to develop a telemedicine application based on a scaleable architecture to be used in Eastern Europe in the field of cardiology. More information is available on the Internet: http://samta.offis.uni-oldenburg.de/

2. Non-Pixel Data in DICOM

Even though the large amounts of image data promise the best compression results compared to other data also stored in DICOM data sets, we are concentrating on the non-pixel data. The reason for this is that with the emergence of Structured Reporting and Presentation States, image and related information (e. g. a report) can be transmitted separately. In the original document [5] this paper is based on, innovative image compression techniques are also examined.

The basic data structure in DICOM is the tag based attribute. A data set (e. g. a CT image) consists of a set of attributes, each of them identified by a unique numerical pair (group and element number). Furthermore the attribute contains the number of bytes used for its value and the value itself. The format of the stored value depends on the Value Representation (VR). DICOM currently defines 26 of these data types (e. g. text string or integer number). It is optional to transmit information about the VR since the tag already uniquely identifies the attribute and for most attributes only one value representation is allowed (and defined in a central data dictionary where it can be looked up). DICOM also provides a mechanism to specify complex data structures as data sets. The so-called Sequence of Items contains a number of data sets, each of them capable of also containing other sequences of items. The length of each item can be specified explicitly or implicitly using special delimiters.

3. Suitability for Narrow-Band Networks

To assess DICOM's suitability for narrow-band networks in its present state, we concentrate on the data structure level first since the network services are based on the same structures, namely data sets. When transmitting DICOM data via network, there is only a little overhead for the TCP/IP and DICOM Upper Layer protocol as well as the DICOM command messages. Only if compatibility with other DICOM devices is not required and the bandwidth is extremely low it may be worthwhile to choose completely different transmission schemes.

The biggest resource problem in the transmission of DICOM images is not the protocol overhead but the large amount of data. Up to now only the compression of pixel data is defined in the standard. The compression of other data (header information) has been disregarded in the past. One reason might be that the pixel data makes up more than 95% of the overall data for most data sets. On the other hand, new modalities introduce more and more of non-pixel data.

Redundant information like value representation, group length and item delimitation can be omitted, but this may only cause significant space savings when transferring data sets without large attributes (like pixel data). Data stored in private attributes may also be a problem when trying to reduce the total size of a data set. If the content or even the value representation of an attribute is unknown to the transmitting application, any reduction becomes difficult.

Many attributes are stored in textual form. There are lots of value representations defined in the DICOM standard handling with textual information. Though it is not the most efficient way to store, e. g. integer numbers, the data exchange between different platforms using different binary formats to store integer numbers is facilitated. However, the result is much redundancy in the DICOM data set and, therefore, a waste of storage space.

Binary data makes up the vast majority of the data set size. Look-up tables can grow up to 128 Kbyte and overlays can be as large as image pixel data. For instance colour look-up tables can be split into several segments defining a linear or non-linear transform or referencing other segments. Furthermore, the bits of an overlay plane can be merged into the unused bits of the pixel data attribute.

4. Experiments

Some of the considerations described above have been tested in practice to see whether there is a potential for a better compression than available nowadays. The experiments have been performed using a large set of medical image data covering all modalities used in cardiology and a number of modalities used outside this medical field. The DICOM images for our test set have been taken from publicly available resources (demo CDs and web sites) or created with our own tools if necessary. The test set covers the most interesting DICOM structures mainly having the compression of non-pixel data in mind.

The test set used for the experiments described in the following consists of:

- medical images from the prevalent modalities in cardiology: X-ray angiography (XA), ultrasound (US);
- medical images from tomographic modalities: computed tomography (CT), magnetic resonance (MR);
- medical images from other modalities: computed radiography (CR), digital X-ray (DX), radio fluoroscopy (RF), radio therapy (RT), visible light (VL);
- new "non-pixel" objects: presentation state (PR), waveform (WV);
- special DICOM structures: DICOM directory (DICOMDIR), look-up table (LUT), overlay plane (OVL), curve (CRV);
- DICOM messages: a large amount of small DIMSE messages acquired from a typical query/retrieve communication.

For the compression of binary, non-pixel data contained in DICOM objects we have decided to use the ZIP tool implementing the Z-lib [7] algorithm. ZIP uses the LZ77 algorithm with hashing plus a secondary static Huffman coding on a block basis and is available for most operating systems. Apart from its public availability the main reason to choose ZIP was that the Z-lib algorithm is supported by the SSH [8] protocol and, therefore, allows to transfer the results obtained for file compression to network transmission.

Since the compression of pixel data is not investigated here, we first have removed the pixel data element from all objects in our test set. Then we have compressed the results with the ZIP tool using different levels of compression. In addition to the optimal compression ratio we have also tried "minimum compression" which is the fastest method. In this phase of our experiments we have not measured the time used for the compression and decompression, but when transferring data via network, the overall transmission time (including compression / decompression) is very important.

Table 1 shows the results of this part of our experiments. The main conclusions are:

- Modalities or structures which can be compressed very well (more than 60% after the possible removal of the pixel data) are e. g. RT, DICOMDIR and OVL since they contain many textually encoded information (14.4 Kbyte on average for our RT test images).
- XA and US which are more interesting for cardiology can be compressed at least by 48% to 58% on average. Compared to the original image size, though, the ratio decreases to less than 0.01% for XA.
- In some cases the compression ratio strongly depends on the particular object. For PR e. g. the best result was 50% compared to 37% on average.
- One idea which has also been tested was to combine several objects (e. g. sharing common information) before compressing. Compressing the header information of different CT and MR series, the compression ratio could be increased by factor 2 on average.

- Re-ordering of the attribute components (group number, element number, VR, value, length) can enhance but also worsen the compression ratio. The best order always depends on the specific object.
- Overlay planes are not very common in real-world DICOM objects but they can be compressed very well (more than 80%) since they typically consist of large uniform areas. With JBIG [6] compression, the ratio can even be enhanced up to 98%.
- The difference between best (ZIP-9) and worst (ZIP-1) compression method is not very large when testing a whole set of images. This compression value influences the size of the "window" in which the compression algorithm searches for redundancy. The default window size produces almost the same results but is significantly faster.

modality or structure	average size		portion of	compression ratio using method			
	of original object	of object without pixel data	non-pixel data vs. original size	ZIP-1 vs. object w/o pixel data	ZIP-9 vs. object w/o pixel data	ZIP-ordered vs. object w/o pixel	ZIP-ordered vs. original size
XA	22375 K	1.4 K	< 0.1%	45.0%	45.8%	48.0%	< 0.01%
US	577 K	2.0 K	0.3%	54.2%	55.6%	58.0%	0.20%
CT	515 K	3.5 K	0.7%	46.4%	49.0%	51.7%	0.36%
MR	514 K	2.0 K	0.4%	35.5%	36.6%	38.8%	0.15%
CR	7223 K	1.4 K	0.0%	35.0%	35.8%	38.1%	0.01%
DX	8699 K	5.5 K	0.1%	22.1%	22.9%	23.1%	0.01%
RF	1025 K	1.5 K	0.1%	39.5%	40.2%	43.2%	0.06%
RT	441 K	14.4 K	3.3%	74.6%	77.6%	75.2%	2.46%
VL	1286 K	0.8 K	0.1%	27.8%	28.2%	29.9%	0.02%
PR	1.5 K	1.5 K	100.0%	35.3%	36.9%	37.6%	37.56%
WV	69.7 K	69.7 K	100.0%	46.7%	49.3%	46.7%	46.74%
D-DIR	243 K	243 K	100.0%	63.9%	65.2%	65.5%	65.46%
LUT	4.7 K	4.7 K	100.0%	24.4%	25.8%	24.6%	24.58%
OVL	334 K	14.9 K	4.5%	82.4%	84.0%	83.3%	3.71%
CRV	257 K	1.2 K	0.5%	32.4%	33.1%	36.3%	0.17%
DIMSE	0.1 K	0.1 K	100.0%	-	-	-	-

Table 1: Results of the experiments compressing non-pixel data

5. Recommendations

Based on the experience gained during the experiments we have extracted a short list of implementable recommendations. We believe that the following four extensions promise best improvements with justifiable effort:

- A transparent "compression layer" should be integrated into the protocol stack to reduce the total amount of data transferred via the narrow-band network. Especially DICOM objects with large non-pixel portions would profit from it. Currently SSH would be a good "ready-to-use" solution because the integration of a compression algorithm is already specified with the Z-lib support. In the longer term TLS [9] should be used instead because of its flexibility and standardization.

- Currently DICOM specifies no compression for overlay plane data at all. We propose to introduce a JBIG compression similar to the JPEG compression of pixel data which is already specified in the standard. Of course, the improvements achieved by this extension strongly depend on the use of overlay planes in real-world DICOM objects.
- Especially for the use of unstable network connections and the transmission of very large DICOM objects we recommend to introduce a "re-get" functionality as part of an extended storage and query/retrieve service. After a break-down the transmission of the last object would not need to be restarted but could be resumed with the last not completely transferred data packet instead.
- Finally, the compression ratio of the algorithms used and the effort put into compression (e. g. none, only on transport layer, additional measures like JBIG or a re-ordering of data elements) should be adaptable automatically. This could be done by introducing a means of monitoring the current overall time needed for compression, transmission and decompression. In particular for connections with changing bandwidths it is important that the available bandwidth is used as well as possible.

6. Conclusion

We have found it possible to reduce the size of header information in DICOM data sets significantly. However, the contribution of a header compression to the overall compression ratio (including compressed pixel data) is very small. Depending on the data to be compressed, a "manual" optimisation can improve or worsen the result of the following compression on the network transmission layer. Our conclusion is that compression of DICOM header information is only worthwhile for data sets which mainly consist of non-pixel data. Since new DICOM modalities introduce more and more of such data, the importance of this issue is likely to increase. In addition, many DICOM network services are based on "textual" data structures. In particular for use with narrow-band networks, extensions as proposed in this paper could be a solution to save valuable bandwidth..

7. References

1. NEMA Standards Publication PS 3.x, *Digital Imaging and Communications in Medicine (DICOM)*, National Electrical Manufacturers Association, 2101 L Street, N. W., Washington, D. C. 20037, 1992-99.
2. DICOM Standards Committee, Working Group 8 - Structured Reporting: Digital Imaging and Communications in Medicine (DICOM), *Supplement 23: Structured Reporting Storage SOP Classes*, Letter Ballot, October 1999.
3. DICOM Standards Committee, Working Group 11 - Display Function Standard: Digital Imaging and Communications in Medicine (DICOM), *Supplement 33: Grayscale Softcopy Presentation State Storage*, Final Text, September 1999.
4. DICOM Standards Committee, Working Group 1 - Cardiac and Vascular Information: Digital Imaging and Communications in Medicine (DICOM), *Supplement 30: Waveform Interchange*, Letter Ballot, November 1999.
5. J. Riesmeier, M. Eichelberg, J. Punys, V. Punys: *Recommendations for an extension of MEDICOM/ DICOM for narrow-band networks*, Deliverable D06.1, INCO-COPERNICUS Project Nr. PL961144 - SAMTA, January 2000.
6. ITU-T Rec.T.85: *Application profile for Recommendation T.82 - Progressive bi-level image compression (JBIG coding scheme) for facsimile apparatus*, 1988.
7. P. Deutsch, J. L. Gailly: *ZLIB Compressed Data Format Specification version 3.3*, RFC 1950, 1996.
8. T. Ylonen et al.: *SSH Protocol Architecture / SSH Transport Layer Protocol*, Internet draft, 1999.
9. T. Dierks, C. Allen: *The TLS Protocol - Version 1.0*, RFC 2246, The Internet Society, 1999.