

---

# Alternative Memory Models for ITK Images

*Release 1.02*

Dan Mueller<sup>1</sup>

November 14, 2009

<sup>1</sup>Philips Informatics Infrastructure (PII), Philips Healthcare, Best, Netherlands

## Abstract

By default ITK images employ a contiguous memory model. This means pixel elements are stored in a single 1-D array, where each element is adjacent in memory to the previous element. However in some situations a contiguous memory model is not desirable. This document describes three alternative memory models: slice contiguous, sparse, and single-bit binary images. A slice contiguous image is a three-dimensional image whereby each slice is stored in a contiguous 1-D array, but the slices are not necessarily adjacent in memory. Slice contiguous images are well suited for interoperability with applications representing images using DICOM. A sparse image is an  $n$ -dimensional image in which each pixel is stored in a hash table data structure. This memory model is well suited for images with very large dimensions, but few pixels that are actually relevant. A single-bit binary image is an  $n$ -dimensional image that internally stores a boolean as a single-bit, rather than the typical eight-bits. Single-bit images allow very compact representations for on-off binary masks. Source code, tests, and examples are provided to allow easy reproduction and use.

Latest version available at the [Insight Journal](http://hdl.handle.net/10380/3068) [ <http://hdl.handle.net/10380/3068> ]

Distributed under [Creative Commons Attribution License](#)

## Contents

<b>1</b>	<b>Introduction</b>	<b>2</b>
<b>2</b>	<b>Proposed Image Memory Models</b>	<b>2</b>
2.1	Slice Contiguous Images	2
2.2	Sparse Images	4
2.3	Single-bit Binary Images	4
<b>3</b>	<b>Performance</b>	<b>5</b>
<b>4</b>	<b>Conclusion</b>	<b>6</b>
<b>5</b>	<b>Acknowledgments</b>	<b>6</b>

---

## 1 Introduction

ITK images currently employ an implicit contiguous memory model. By implicit we mean that the memory model is tightly coupled with many other classes within the toolkit. By contiguous we mean that pixel elements are stored as a 1-D array, where each element is adjacent in memory to the previous element. This underlying pixel array is not hidden behind a layer of abstraction but can be easily accessed using for example `itk::Image::GetBufferPointer()`.

Performing a quick grep over the codebase reveals numerous classes which directly access the image pixel array in this manner. Such tight coupling makes it difficult to abstract away the underlying memory model used by the image to store the pixel data. Some of the important classes which directly access the pixel array using `itk::Image::GetBufferPointer()` include: image iterators, image adaptors, image file readers/writers, and the VTK image export filter. A number of other minor classes also directly access the pixel array, such as: octree, watershed segmenter, BSpline deformable transform, optimized Mattes mutual information metric, etc.

The existence of classes which directly access the pixel array, as well as ITK's strict backward compatibility policy, makes it difficult — but not impossible — to introduce images with alternative memory models. This paper uses the same approach employed by `itk::VectorImage` to realize three new memory models for ITK images. The new image types do not require changes to the toolkit, and function with a majority of the existing filters and iterators. However they are only interoperable with classes which do not directly access the pixel array. For example it is not possible to use image IO, export to VTK, or use the optimized registration framework with the proposed images.

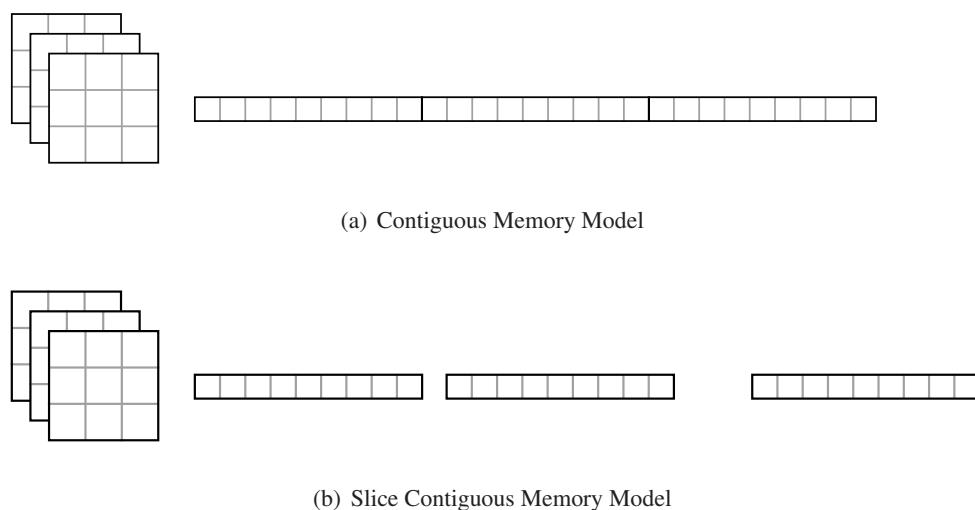


Figure 1: Contiguous and slice contiguous memory models.

## 2 Proposed Image Memory Models

### 2.1 Slice Contiguous Images

A slice contiguous image is implicitly three-dimensional. Each slice is stored in a contiguous 1-D array, but the slices are not necessarily adjacent in memory. This representation is shown in Figure 1. The new `itk::SliceContiguousImage` may be useful for interoperability with existing code representing images

using a slice contiguous memory model. This scenario is common for third party applications using DICOM.

This paper proposes an implementation for an image with such a memory model. This is realized by creating a new image class with custom pixel and neighborhood accessor functions; similar to `itk::VectorImage`. The pixel and neighborhood accessor functions provide a layer of indirection for iterators to access the image pixel data. Given the incoming offset for a contiguous memory model, the new accessors compute the index of the slice containing the pixel and the offset within that slice. The new image class is templated over the pixel type; it is not templated over the number of dimensions because it is always three.

One important difference between `itk::Image` and `itk::SliceContiguousImage` is that for the later `GetBufferPointer()` always returns NULL ie. the pixel buffer makes no sense. Ideally this method would not even exist for slice contiguous images, but unfortunately too many other classes assume its existence.

A slice contiguous image is created in a very similar fashion to a “normal” image:

```

1 // Typedefs
2 typedef float PixelType;
3 typedef itk::SliceContiguousImage< PixelType > SliceContiguousImageType;
4
5 // Create image
6 SliceContiguousImageType::Pointer image = SliceContiguousImageType::New();
7 SliceContiguousImageType::IndexType start;
8 SliceContiguousImageType::SizeType size;
9 start.Fill( 0 );
10 size[0] = sizeSlice[0]; size[1] = sizeSlice[1]; size[2] = 4;
11 SliceContiguousImageType::RegionType region( start, size );
12 image->SetRegions( region );
13 image->Allocate( );

```

An additional step requires the pixel container to be configured with a list of slice pointers:

```

1 // Set slice pointers
2 SliceContiguousImageType::PixelContainer::SliceArrayType slices;
3 slices.push_back( slice2->GetBufferPointer() );
4 slices.push_back( slice4->GetBufferPointer() );
5 slices.push_back( slice3->GetBufferPointer() );
6 slices.push_back( slice1->GetBufferPointer() );
7 SliceContiguousImageType::PixelContainerPointer container =
8     SliceContiguousImageType::PixelContainer::New();
9 container->SetImportPointersForSlices( slices, size[0]*size[1], false );
10 image->SetPixelContainer( container );

```

The slice contiguous image can now be used like any other image:

```

1 // Use the image with a filter
2 typedef itk::ShiftScaleImageFilter<SliceContiguousImageType,
3     SliceContiguousImageType> ShiftScaleFilterType;
4 ShiftScaleFilterType::Pointer filter = ShiftScaleFilterType::New();
5 filter->SetInput( image );
6 filter->SetShift( 10 );
7 filter->Update( );

```

## 2.2 Sparse Images

A sparse image is an  $n$ -dimensional image in which each pixel is stored in a hash table data structure. Each time a pixel is set, a new offset-value pair is added to the hash table. Such a memory model means that little or no memory is allocated when the image is created, but the memory footprint grows as more and more pixels are set. This memory model is well suited for images with very large dimensions, but few pixels which are actually relevant. It should be noted that (at least for the moment) filters using images with this memory model must be single threaded.

A sparse image can be created in the typical fashion:

```
1 // Typedefs
2 const unsigned int Dimension = 2;
3 typedef unsigned short PixelType;
4 typedef itk::SparseImage< PixelType, Dimension > SparseImageType;
5
6 // Create image
7 SparseImageType::Pointer image = SparseImageType::New();
8 SparseImageType::IndexType start;
9 start.Fill( 0 );
10 SparseImageType::SizeType size;
11 size.Fill( 1000000 ); // try this with a "normal" image!
12 SparseImageType::RegionType region( start, size );
13 image->SetRegions( region );
14 image->Allocate( );
```

A sparse image requires a “background” value to be specified for undefined pixels:

```
1 image->FillBuffer( 100 );
```

Pixels can be retrieved or set directly (using `Get/SetPixel`) or via an iterator:

```
1 image->FillBuffer( 100 );
2 SparseImageType::IndexType indexA = {100, 100};
3 PixelType pixelA = image->GetPixel( indexA );
4 SparseImageType::IndexType indexB = {10000, 10000};
5 image->SetPixel( indexB, 5 );
6 PixelType pixelB = image->GetPixel( indexB );
```

## 2.3 Single-bit Binary Images

It is very common in image processing to create mask images which indicate regions which are “on” or “off”. Currently the smallest pixel element available in ITK for representing such images is `unsigned char` which is stored in a single byte (eight-bits). Even though a `bool` value can only represent 0 or 1, it too is stored in a single byte (you can check this by writing out `sizeof(bool)`).

Single-bit binary images internally represent each pixel as a single-bit. As such the memory footprint for on-off masks can be lowered by (nearly) a factor of eight. Similar to slice contiguous images, single-bit binary images provide custom pixel accessor functions which convert the incoming offset to the relevant bit mask for the underlying data storage. Unlike slice contiguous images, single-bit binary images fit slightly better within the existing ITK framework so `GetBufferPointer()` makes sense in this context.

A single-bit binary image can be created in the typical fashion:

```

1 // Typedefs
2 const unsigned int Dimension = 2;
3 typedef bool PixelType;
4 typedef itk::SingleBitBinaryImage< Dimension > BinaryImageType;
5
6 // Create image
7 BinaryImageType::Pointer image = BinaryImageType::New();
8 BinaryImageType::IndexType start;
9 start.Fill( 0 );
10 BinaryImageType::SizeType size;
11 size.Fill( 65 );
12 BinaryImageType::RegionType region( start, size );
13 image->SetRegions( region );
14 image->Allocate();
15 image->FillBuffer( true );

```

The bits are stored in blocks of 32, so the above code will actually allocate a buffer with size  $96 \times 96$ . As with all the presented image class, the binary image can be used with any iterator and/or filter which does not directly access the pixel buffer via `GetBufferPointer()`.

### 3 Performance

Although the memory models presented above may be advantageous for reducing memory usage in certain scenarios, they have a performance penalty. Accessing pixels stored in a contiguous array can be highly efficient, whereas the three new images require additional computation each time a pixel is accessed. This section provides a simple analysis of the performance for each of the proposed images.

The performance test measured four properties: (1) time to allocate the buffer, (2) time to fill the buffer, (3) time to set all pixels using an iterator, and (4) time to get all pixels using an iterator. Each test was run on a  $256 \times 256 \times 256$  image, executed on my notebook (Intel Core 2 Duo, T7250 @ 2GHz, 3GB RAM, Windows Vista SP1 32-bit) a total of 5 times with mean times (in seconds) reported in the table below.

As can be seen, all the images have similar values for buffer allocation and getting pixels using an iterator. The sparse image was the fastest for filling the buffer because no memory is actually set at this moment, only a single “background” value. However it was also the slowest (by quite a margin) for setting pixels using an iterator, because each pixel being set must be added to the hash table. The single-bit binary image was comparably fast for filling the buffer because the pixels are set in groups of 32-bits, rather than individual elements.

Image Type	Allocation	FillBuffer	Iterate Set	Iterate Get
Contiguous	< 0.01	0.05	0.05	0.02
Slice Contiguous	0.01	0.05	0.13	0.02
Sparse	< 0.01	< 0.01	6.53	0.02
Binary	< 0.01	< 0.01	0.07	0.02

Table 1: Performance timings in seconds for a  $256 \times 256 \times 256$  image.

## 4 Conclusion

This submission proposed three images with alternative memory models: slice contiguous, sparse, and single-bit binary images. For the moment there is no IO support for any of the proposed images, however it might be possible to provide such support in the future. The proposed images should work seamlessly with most of the existing ITK filters — assuming they access the pixel data using iterators rather than `GetBufferPointer()`. These new images are just a few of the possible alternate memory models and should provide a good guide for others wanting to develop their own.

## 5 Acknowledgments

Thanks to Karthik Krishnan (Kitware) and Glen Lehmann (Calgary Scientific) for helpful discussions.