

MtxVec v1.5

A programming guide to MtxVec

1 Introduction

MtxVec is an object oriented numerical library. There have been quite a few attempts to make such a library for numerical analysis object oriented. Many people have tried to force a class hierarchy on numerical algorithms, but the most successful libraries today do not use class hierarchies. For example, the basic element of Matlab (by Mathworks) is one single object. That object type can represent real or complex single value, vector, matrix or sparse matrix or even a string. MtxVec introduces two classes:

TVec for vector operations

TMtx for matrix operations

Both objects can hold either real or complex data. Both classes should be considered as sealed objects. (that is to say - no classes should be derived from them). Both classes hold a list of methods which may be called like this:

```
vector_object_a.Add(vector_object_b);
```

If the method can put its result in one object, then the result is placed in the object on the left. In the following example the result is placed in the objects on the right:

```
a.CplxToReal(Real, Imag);  
a.CartToPolar(Amplt, Phase);
```

Its useful to remember this when mixing TVec and TMtx types. A matrix operation which has TVec type as result will be a part of the TVec class and a vector operation which has a TMtx type result will be a part of TMtx class.

1.1 Indexes, ranges and subranges

Most methods support indexing:

```
a.Copy(b, 2,0,10);
```

'a.Copy' means 'copy' 10 elements of "b" from index 2 to "a" starting at index 0 of a. If there are no index parameters, the size of the target object will be set automatically. An alternative means for indexing is to use SetSubIndex or SetSubRange methods:

```
b.SetSubRange(2,10);  
a.Copy(b);
```

Which is the same as:

```
b.SetSubIndex(2,11);  
a.Copy(b);
```

The use of SetSubRange and SetSubIndex is recommended because it employs memory reuse, which takes advantage of the CPU cache, which in turn improves performance. There are other types of indexing where there is a need to apply an operation to specific non-continuous indexes within a vector or matrix. This can result in heavy performance penalties (heavy means by a factor of 100-300) for some numerical algorithms. The entire CPU architecture is based on the assumption that memory is accessed by consecutive memory locations in 90% of cases. It is therefore best to first gather the scattered data into one dense vector, perform math operations and then scatter the gathered data back to the original location:

```
a.Gather(b,nil,indIncrement ,2);  
a.Log10;
```

```
a.Exp;  
b.Scatter(a,nil,indIncrement,2);
```

This code will copy every second element from b to a, apply math and then scatter the result back to b without affecting other values in b. The Gather and Scatter methods can also accept an index or a mask vector. To access elements of an index vector via IValues array:

```
a.IValues[0] := 1;
```

IValues points to the same memory as Values and CValues and is simply an integer array type pointer. There are more routines that can help with scattered data:

```
a.FindMask(b, '=', c);
```

The method will return ones for all indexes where b and c have a matching value and zeros elsewhere. FindAndGather can be used to find all indexes within b where values are different from NAN (not a number) and apply processing only to those values:

```
a.FindAndGather(b, '<>', NAN, Indexes);  
a.Scale(2);  
b.Offset(1);  
a.Scatter(b, Indexes);
```

1.2 TVec and TMtx methods as functions

Ideally a mathematical expression could be written like this:

```
a := a*b + b;
```

For vectors and matrices in Delphi this can not be done. The closest syntax allowed is this:

```
a.Add(c.Mul(a,b),b); //where a,b,c are TVec objects
```

Almost every method of TVec and TMtx returns Self. This allows nesting of calls like in the example above. Syntax like this will result in a memory leak:

```
a := c.Mul(a,b); //don't do this
```

1.3 Create and destroy

Before objects can be used, they have to be created. TVec and TMtx objects can be created and destroyed in the standard way:

```
a := TVec.Create;  
B := TMtx.Create;  
try  
....  
finally  
    a.Destroy;  
    B.Destroy;  
end;
```

Or in a fast way, by using object cache:

```
CreateIt(a);  
CreateIt(B);  
try
```

```
....  
finally  
    FreeIt(a)  
    FreeIt(B);  
end;
```

Object cache is a set of objects, which are created when the application is started. When a call to CreateIt is made, no object actually gets created. The CreateIt procedure simply assigns a pointer to an already created object to the parameter. That is not all since the already created object has some memory allocated and there is no new memory allocated until some default size is exceeded. This type of memory allocation (call it preallocation) is speedier by a factor of 2 and in some cases even more. It is difficult to predict the actual effect on the entire application, where the gains could be significant. The use of object cache is not significant only because the calls to Create/Destroy and GetMem/FreeMem are never made, but also because it increases memory reuse.

1.4 Complex data

Both TVec and TMtx can hold real and complex data. Here is an example:

```
a.Length = 10;  
a.Complex = True;
```

a.Length now becomes 5. Setting the complex property will simply halve or double the length property of the vector. The allocated memory will not change. There is a need however to view that memory as a real or as a complex array:

```
a.Values[0] := 1;  
a.CValues[0] := Cplx(2,3);
```

a.Values[0] now becomes 2, because both Values and CValues arrays point to the same memory. The only difference between them is that one is of type double and the other is of type TCplx (TCplx = record Re,Im: double; end;).

1.5 MtxVec types

MtxVec declares many different types, but some should be mentioned explicitly:

1.5.1 TSample

This type is used everywhere where it is necessary to declare a floating point number. It can be declared as double or single (in Math387).

```
{ $IFDEF TDOUBLE } type TSample = double; { $ENDIF }  
{ $IFDEF TTSINGLE } type TSample = single; { $ENDIF }
```

By using a DEFINE statement in bdsppdefs.inc file, the precision in which MtxVec runs can be switched between double and single.

1.5.2 TCplx

Declared as:

```
TCplx = packed record
```

```
Re, Im: TSample;  
end;
```

This is the default complex number type used by MtxVec. Many object oriented libraries declare the complex type as an object type (TComplex = class(TObject)). This is not very good, because then the consecutive elements within an array are no longer stored at consecutive memory locations. The result are extreme performance penalties. (up to 300x).

1.5.3 TSampleArray, TCplxArray

Declared as:

```
TSampleArray = array of TSample;  
TCplxArray = array of TCplx;
```

2 Programming style

Every programmer has a preferred style of programming: different indentation, different variable naming, different coding style (use of exceptions, for loops, dynamic memory allocation etc.). This section lists the recommended coding style for programming with MtxVec.

2.1 Try-finally blocks.

Every time a call is made to CreateIt/Freelit or Create/Destroy pair, it should be placed within a try-finally block like this:

```
var a,b,c,d: TVec;  
begin  
    CreateIt(a,b,c,d);  
    try  
        ..  
        //Your code here.  
    finally  
        FreeIt(a,b,c,d);  
    end;  
end;
```

These has two purposes:

If there is an exception within the try-finally block, the allocated objects and memory will be freed and the program user will be able to retry the calculation with other parameters. It is now easier to track what is created and what is destroyed, because it is clearly visible where create and where destroy is called. MtxVec has internal variables tracking the state of the object cache. If those variables are not zero when application terminates, a call to Freelit has somewhere been left out.

Do not write code like this:

```
CreateIt(b);  
  
//...some code here  
  
CreateIt(a);  
yourProc(a,b);  
Freeit(b);  
  
//.. some code here..  
FreeIt(a);
```

This makes it difficult to see, if all calls to CreateIt have FreeIt pairs. It is also a good rule of housekeeping to group the code allocating the memory separately from the code doing calculations. This makes code much more readable.

2.2 Raising exceptions

Because all code is now protected with try-finally blocks, exceptions can now be raised safely to indicate an invalid condition. When the user tries to perform calculation with an MtxVec application, this is what will happen:

- 1.) Allocate memory for the calculation.
- 2.) Start calculation
- 3.) Display results.
- 4.) Free allocated memory and resources.

If during the calculation an error condition is encountered, because the data is not valid, raising an exception will first Free allocated memory and resources and then display a message box stating what the error was. It is important that memory was freed, because now the user can retry the calculation with new data or parameters, without the need to restart the application to reclaim the lost memory.

2.2.1 Invalid parameter passed:

```
begin
  if a = nil then raise Exception.Create('a= nil');...
  ...
end;
```

This will pass an exception to the higher-level procedures, which will free any allocated memory and exit. Once the exception reaches the highest-level routine a message box will be displayed with text "a = nil" and an OK button.

2.2.2 Reformat the exception

Once an exception has been caught, one might want to notify the user, not of some variable being nil, but actually which procedure failed:

```
try
  ...
except
  on E: Exception do
    raise Exception.Create(YourMessageHere + ' : ' E.Message);
end;
```

This example retains the old messages, adds custom string to it and then raises the exception again, this time with a new message. Every exception will show up in the program as a message box. ShowMessage or MessageDlg or MessageBox should not be used to indicate an error condition. An exception should be raised instead. Raising an exception will safely exit all nested routine calls, free associated memory and finally also show the message box.

2.3 By all means indent code.

A procedure should look like this:

```
var a,b,c,d: TVec;
    i: integer;
```

```
begin
  CreateIt(a,b,c,d);
  try
    //   SomeCodeHere..
    for i := 0 to a.Length-1 do
      begin
        //   MoreCodeHere.....
      end;
    finally
      FreeIt(a,b,c,d);
    end;
end;
```

2.4 Do not create objects within procedures and return them as result:

```
procedure GetVector(var a: TVec);
begin
  CreateIt(a);
end;
```

This makes it difficult to track CreateIt/Freelit and Create/Destroy pairs and breaks a few good rules of programming.

2.5 Use CreateIt/Freelit only for dynamically allocated objects whose lifetime is limited only to the procedure being executed.

All objects created within a routine should be destroyed within that same routine. If TVec or TMtx are global objects, make them a part of an object or component. Global objects are those, which are not created and destroy very often and might persist in memory throughout the life of an application. This rule should be followed in order not to waste the object cache. The purpose of object cache is to allow speedy memory allocation and deallocation. Where this is not needed, it should not be used, because that could slow down other routines using it:

- 1.) Object cache might run out of precreated objects and calls to CreateIt/Freelit would result in direct calls to Create/Destroy.
- 2.) Object cache size would have to be increased to prevent (1) and the entire application would require more memory.

3 Accessing values of TVec and TMtx

3.1 Default array property access.

Example:

```
var a: TMtx;
begin
  CreateIt(a)
  try
    a[1,0] := 2;
  finally
    FreeIt(a);
  end;
end;
```

Default array property allows a clean access, but is not very fast. It's use for large matrices is discouraged. The default array property performs explicit range checking.

3.2 Dynamic array pointer.

Example: `a.Values[1,0] := 2;` {where a is TMtx object}

Values is a dynamic array pointer. The access is significantly faster then with default array properties.

3.3 Direct dynamic array pointer.

Example:

```
var ap: T2DSampleArray;
begin
    CreateIt(a);
    try
        a.Length := 4;
        Enlist(a,ap); {is the same as: ap := a.Values}
        ap[1,0] := 2;
        Dismiss(ap)
    finally
        FreeIt(a)
    end;
end;
```

This is the fastest method and recommended for long loops and large matrices. The speed of access is equivalent to static arrays.

Handling complex data :

Default array property access is not available for complex data because the object can have only one default array property. The following access methods are semantically equivalent:

`a.CValues[1,0] := Cplx(2,0);` {Cplx is a function that returns a TCplx record type}

...

```
var ac: T2DCplxArray;
begin
    CreateIt(a); {similar as a := TVec.Create;}
    try
        a.Rows := 4;
        a.Cols := 4;
        a.Complex := True;
        Enlist(a,ac) {is similar to: ac:= a.CValues}
        ac[1,0] := Cplx(2,0);
        Dismiss(ac);
        //.... {same as:}
        a.Values[2,0] := 2;
        a.Values[3,0] := 0;
        //.... {same as:}
        a.CValues[1,0].Re := 2;
        a.CValues[1,0].Im := 0;
        ...
    finally
        FreeIt(a); {similar as: a.Destroy; a := nil}
    end;
end;
```

There is one important consideration, because of the way how Delphi works with dynamic arrays. This should not be attempted:

```
var ar: TSampleArray;
    a: TVec;
begin
    CreateIt(a);
    try
        a.Length = 10;
        a.SetSubIndex(2,2);
        ar := a.Values; //problem here
        ar[0] := 1;
    finally
        FreeIt(a);
    end;
end;
```

a.Values should **not** be assigned to TSampleArray variable after a call to SetSubIndex or SetSubRange. This could corrupt the data in the "a" vector. Enlist/Dismiss pair should be used instead.

4 Memory management

The memory for TMtx is allocated by setting the Rows and Cols properties:

```
a := TMtx.Create;
a.Rows := 4; {allocates nothing}
a.Cols := 4; {a now holds 16 elements}

a.Rows := 0; {deallocates memory}
a.Destroy;   {same as a.Rows := 0; a.Cols := 0; a.Destroy;}

CreateIt(a); //similar to a := TMtx.Create;
a.Size(4,4,False); //same as: a.Complex := False; a.Rows := 4; a.Cols := 4;
FreeIt(a); // simiarl as a.Destroy; a := nil;
```

The complex property should be set before setting the Cols property. All arrays are zero based. (The first elements is always at index 0).

There are some special issues that need to be taken in to account when working matrices. TMtx interfaces highly optimized FORTRAN code. Dynamic memory allocation of two dimensional matrices in Delphi is not done in one single block as expected by fortran routines. For that purpose, TMtx uses it's own memory allocation to dynamically allocate two dimensional arrays in a single block of memory. Therefore, there are two more pointers available from TMtx:

```
a.Values1D[i]
a.CValues1D[i]
```

These two pointers point to the same memory location as Values and CValues pointers. But instead of accessing the elements by rows and columns, they see the whole matrix as one-dimensional array. To access matrix elements:

```
a1 := a.Values1D[i*Cols+j];
```

This will access the same matrix element as:

```
a1 := a.Values[i,j];
```

The preferred method for memory allocation, is by using the Size method:

```
a.Size(4,4,False);
```

Size method will ensure that no more memory is allocated than necessary when resizing. Imagine a 5x10000 matrix, being resized to 10000x5, but the rows are set to 10000 first creating a matrix with 10000x10000 elements, possibly causing an out of memory message.

Matrix data is stored in row-major ordering of C and PASCAL and not in column major ordering of the FORTRAN language. All appropriate mappings are handled internally.

5 Range checking

Due to dynamic memory allocation in one single block, the array range checking is not performed by the compiler for matrices, but TMtx does perform explicit range checking, if the default array property is used:

```
a[i,j] := 2; {performs the following:
```

```
if i > a.Rows-1 then raise Exception...
if i < 0 then raise Exception...
```

```
if j > a.Cols-1 then raise Exception...
if j < 0 then raise Exception...
```

```
a.Values[i,j] := 2}
```

If there is ever a question about how much memory to allocate, or if some algorithm might be crossing the bounds of arrays, default array property can be used. Once the code is debugged, it can be switch back to high speed array pointers.

In-place/not-in-place operations

As opposed to TVec, TMtx does not strictly separate between in-place/not-in-place operations. The reason is the high computational complexity of most matrix operations, thus making that feature almost obsolete. In case of very large matrices the memory requirements would become a problem (10 000 x 10 000 matrix requires 800MB storage). In such cases the user should use LAPACK routines directly by adding nmkl to the uses clause. Whenever possible LAPACK performs matrix operations in-place. Often the matrix size can be greatly reduced by using banded matrix format or sparse matrices.

6 Why and how NAN and INF

NAN is short for Not a Number and INF is short for infinity. By default Delphi will raise an exception when a division by zero occurs or an invalid floating operation is performed. By including Math387 unit in the uses clause, the floating point exceptions are automatically disabled. This allows code in loops without the checks, if the function input parameters are within the definition area of that function. This alone speeds up the code, because most of the try-except and if-then clauses used for that purpose can be left out. Instead, you can concentrate on the code itself and let the CPU work out the details. If a division by zero occurs and floating point exceptions are off, then the FPU (floating point unit) will return INF (for infinity). If divide zero by zero is attempted, the FPU will return a NAN (not a number). When working with arrays, this can be very helpful, because the code will not break when the algorithm encounters an invalid parameter combination. It is not until the results are displayed in the table or drawn on the chart that the user will notice that there were some invalid floating point combinations. It might also happen that INF values will be passed to a formula like this: number/INF (= 0) and the final result will be a valid number.

MtxVec offers specialized routines for string to number and number to string conversions (StrToVal, StrToCplx, FormatCplx, FormatSample, StrToSample, SampleToStr) and drawing routines (DrawValues, DrawIt) capable of handling NAN and INF values. By using those routines (located in math387 unit), the user will avoid most of the problems when working with NAN and INF values. StrToSample for example will convert a NAN or INF string to its floating point presentation:

```
var a: TSample;
begin
    a := StrToSample('NAN');
    if IsNan(a) then raise Exception.Create('a = NAN');
end;
```

StrToFloat routine would raise an exception on its own. To test for a NAN and INF value, the first attempt would look like this:

```
if a = NAN then ...
```

This however will not work. NAN and INF are not values which are defined with all the bits of a floating point variable. There are just a few bits that need to be set within a floating point variable which will make it a NAN or an INF. The proper way to test for a NAN and INF are therefore these:

```
if IsNan(a) then ...
if IsInf(a) then ...
if IsNanInf(a) then ...
```

MtxVec methods and routines correctly handle NAN and INF. It is therefore acceptable to write something like this:

```
if a.Find(Nan) > 0 then ..//test if "a" vector holds a NAN value
```

6.1 Delphi 4 and Delphi 5 issues with NAN and INF

Delphi 6 brought one important change when dealing with NAN and INF. In earlier versions of Delphi there were three routines causing problems when floating point exceptions were turned off: Integer(), Trunc(), and Frac(). These routines changed the exception mask of the FPU and re-enabled floating point exceptions. Each time a NAN or INF value was computed by the FPU (floating point processing unit) an exception flag was set, but no exception was raised, because the exceptions were off. When however the floating point exceptions would get enabled, this set exception flags would immediately raise one of the two exceptions: Invalid floating operation (NAN was a result at least once in the previous calculations) or Floating point overflow operation (INF was a result at least once in previous calculations). The only way to resolve this problem is to reset the exception flags of the FPU by calling ClearFPU routine, before a call to Trunc, Frac and Integer. This routine is located in the Math387 unit. Instead of writing code like this:

```
a := Trunc(b); //invalid floating operation might get raised here.
```

one should write it like this:

```
ClearFPU;
a := Trunc(b);
```

One important thing to remember with Delphi 4 and Delphi 5 is to always call ClearFPU before calling any code that might potentially use Trunc, Frac or Integer routines. The most common place for the ClearFPU is right before the values are passed to TeeChart for drawing. (TeeChart uses Trunc). If DrawValues routine from MtxVecTee unit is used, the ClearFPU is called automatically and all NAN and INF values are filtered out from the array and only valid values get drawn. The problem of Delphi 4 and 5 is also present in C++Builder versions 4 and 5.

7 Serializing and streaming

7.1 Streaming with TMtxComponent

All components should be derived from a common ancestor: TMtxComponent. (Declared in MtxBaseComp.pas). This component features the following methods:

- SaveToStream
- LoadFromStream
- SaveToFile
- LoadFromFile
- Assign
- AssignTemplate
- LoadTemplateFromStream
- SaveTemplateToStream
- LoadTemplateFromFile
- SaveTemplateToFile

Most Delphi/CBuilder developers know the first five routines. What is interesting about TMtxComponent is that all components derived from it have all their published properties streamed, without the need to make any changes to the five routines. Therefore, all components derived from TMtxComponent have the capability to store their states (properties) to the stream, file or to assign from another object of the same type. The default component streaming mechanism has a bug when it comes to streaming properties with a declared default value. This has been fixed for TMtxComponent.

The "template" routines are a bit more advanced. They set a special protected property named BlockAssign to True. Property setter routines can then prevent properties to be changed. This is very useful when there is a need to save only "parameters" and not the "data" of the component. The parameters will remain the same, while the "data" will be different next time and there is no point in wasting disk space by saving it

7.2 Streaming of TVec and TMtx.

Both TVec and TMtx type objects will be streamed when declared as published properties of a component. The component must be at least of type TPersistent to support streaming. TVec and TMtx also have their own methods for streaming:

- SaveToStream
- LoadFromStream
- SaveToFile
- LoadFromFile
- Assign

These routines will process all the published properties and data of the TVec and TMtx objects. These methods only save and load data in the binary format. (MtxVec also supports Matrix Market text file format.) However, sometimes it is necessary to read a text file. Here is how this can be done:

7.3 Write TMtx to a text file

```
AMtx.Size(20,20,true);
AMtx.RandUniform(-1,2);
StringList := TStringList.Create;
try
  { use tab = chr(9) as delimiter }
  AMtx.ValuesToStrings(StringList,#9);
  { Save matrix values to txt file }
```

```

    StringList.SaveToFile('ASCIIMtx.txt');
finally
    StringList.Destroy;
end;

```

7.4 Read TMtx from a text file

```

StringList := TStringList.Create;
CreateIt(tmpMtx);
try
    { get matrix values from text file }
    StringList.LoadFromFile('ASCIIMtx.txt');
    { use tab = chr(9) as delimiter }
    tmpMtx.StringsToValues(StringList,#9);
    if CheckBox1.Checked then ViewValues(tmpMtx);
finally
    FreeIt(tmpMtx);
    StringList.Destroy;
end;

```

8 Input-output interface

8.1 Reading and writing raw data

TVec and TMtx have the capability to save their state via `SaveToStream` and `LoadFromStream`. The downside on using these two routines is that it is not possible to save the raw data only. The values of all the properties are always included as the header of the saved data block. Saving raw data only can be achieved by using two other methods: *TVec.WriteValues* and *TVec.ReadValues*. These two methods will read and write to and from TStream descendants only the contents of Values array itself (raw data). When writing to stream, it is also possible to define the precision and consequently the size of the disk space to occupy. Supported precisions include:

```

TPrecision = (prDouble, prSingle, prInteger, prCardinal, prSmallInt, prWord,
prShortInt, prByte, prMuLaw, prALaw, prInt24);

```

MuLaw and ALaw are audio compression standards for compressing 16 bit data to 8 bits. prInt24 is a 24 bit signed integer useful for 24bit digital audio.

When data is being read with `ReadValues`, the type of the data must be explicitly specified. All data types are converted to TSample.

8.2 Formatting floating point values – FloatToString

Math387 unit declares the following routines:

```

function FormatCplx(Z: TCplx; const ReFormat: string = ' 0.###;-0.###';
                    const ImFormat: string = '+0.###i;-0.###i'): string;
function FormatSample(X: TSample; Format: string = '0.###'): string; overload;

```

Both are similar to `FormatFloat` and returned a string representing the floating point number passed as the first parameter. If the `Format` parameter is an empty string, the routines will call the `SampleToStr` and `CplxToStr` routines. They are declared like this:

```

function CplxToStr(const Z: TCplx;const Digits: integer = 0): string; overload;
function SampleToStr(const X: TSample; const Digits: integer = 0;
                    {$IFDEF TTSINGLE} Precision: integer = 7 {$ENDIF}

```

```
{ $IFDEF TDOUBLE } Precision: integer = 15 { $ENDIF }): string; overload;
```

and are similar to FloatToStr. The Digits parameter specifies the minimum number of digits in the exponent. The Precision defines the number of digits to represent the number. Example:

```
a := 12.123456;
myString := SampleToStr(a, 0, 3); // myString = '12.1';

a := 12.123456;
myString := SampleToStr(a); // myString = '12.123456';
```

To convert from string to a floating point number, the following routines can be used:

```
function StrToCplx(const Source: string): TCplx; overload;
function StrToSample(const Source: string): TSample; overload;
```

They are similar to StrToFloat with a few exceptions when it comes to handling NAN and INF values. (See the chapter: "Why and how NAN and INF".)

In general these routines improve the default string-to-number and number-to-string conversions by adding sensitivity to single/double precision and better support for NAN and INF values.

8.3 Printing current values of variables

The most common way to debug an algorithm in the old days was to print the values of the variables to the screen or to the file. Today it is more common to use watches and tooltips to examine the values of variables. In some cases this two approaches do not work because of multithreading. Sometimes it is also desirable to have an algorithm return report on its convergence, like it is the case with optimization algorithms. For cases like this there is a global variable called Report declared in MtxVec.pas unit. Report is of a TMtxVecReport type and is derived from TStringStream. It has been extended with several new methods to support:

- Saving the stream to TStream or a file.
- Write the contents of TVec and TMtx objects to the stream. (as text)
- Specify the text formatting of the floating point values.

Typically the Report object can be used like this:

```
a,b: TVec;
begin
....
Report.Write([a,b],[ 'a', 'b' ]);
Report.NewLine;
Report.SaveToFile('C:\test.txt');
Report.Clear;
```

8.4 Displaying the contents of the TVec and TMtx

8.4.1 As a delimited text

TVec and TMtx have two methods for getting and setting their values to and from a text file. They are called StringsToValues and ValuesToStrings and feature:

- handling of complex and real values.
- accept NAN and INF values
- can get/set only a sub vector or a sub matrix.
- Control the displayed precision

8.4.2 Within a grid.

The best way to display the contents of a matrix is within a grid. TVec and TMtx feature methods declared in MtxDialogs.pas that support getting and setting the contents of the TStringGrid object. These methods are: GridToValues and ValuesToGrid. They feature:

- handling complex and real values
- accept NAN and INF values
- can display column/row headers or not
- can get/set only a sub vector or a sub matrix.
- Control the displayed precision

These methods are also used by the ViewValues method, which is declared in MtxVecEdit.pas.

8.5 Charting and drawing

MtxVec provides two units: MtxVecEdit and MtxVecTee to support the display and charting of the contents of TVec and TMtx. By calling the ViewValues routine a simple editor will be displayed allowing the user to examine the values of a TVec or TMtx object. See Figure 1. This editor can be displayed modally or not. If it is displayed modally, the values can be changed and the contents of the object will also change. If the editor is not displayed modally, the changes will be discarded. If the changes are not to be saved, then the user can freely select the number formatting. If the changes are to be saved, the number formatting must be full precision or otherwise the values will be truncated to the displayed precision. The values can not be edited unless Editable flag from the Options menu is checked. From the editor Chart menu "Series in rows" or "Series in cols" can be selected to draw the displayed values as a chart. Vector or matrix values can also be drawn directly on the chart by calling the DrawIt routine. This routine is located in the MtxVecTee unit. MtxVecTee routine contains a large set of DrawValues routines. This routines copy data from TVec or TMtx to the defined TChartSeries. Adding of new values is optimized for the TeeChart version used and charting can be considerably faster, if DrawValues is used. DrawValues routines also take care of any NAN's and INF's.

```
var a: TVec;
begin
  CreateIt(a);
  try
    a.LoadFromFile('c:\test.vec');
    ViewValues(a); //display a window showing values in "a"
    DrawIt(a); //display a chart of values in "a"
    ...
  finally
    FreeIt(a);
  end;
end;
```

20x20	0-Re	0-Im	1-Re	1-Im	2-Re
0	0.4260	-0.8290	0.1630	0.2030	1.4050
1	1.2040	0.5370	-0.8320	-0.7470	1.9240
2	0.7360	0.5290	1.1640	1.2290	1.3710
3	1.1090	0.4930	0.2380	0.3140	1.3720
4	1.5070	-0.3620	0.4990	-0.0810	0.0880
5	-0.5300	0.8290	0.0600	1.7950	0.6250
6	-0.8650	1.1710	0.9170	-0.5980	-0.7720
7	-0.1710	-0.7680	0.0460	0.9950	-0.2570
8	1.6820	0.2630	-0.3120	0.1320	-0.0260

Figure 1

In the top left corner on Figure 1 is the size of the matrix (in this case 20x20). In the left most column are rows indexed starting with 0. The top row shows column labels. "0-Re" means that this is the first column of a complex matrix and shows the Real part of the complex number, "0-Im" column shows the imaginary component of the complex number stored in the first column of the matrix. On Figure 2 the magnitudes of the values stored in the complex matrix can be seen. The layout of the values is the same as in the matrix editor (the left axis labels should be inverted).

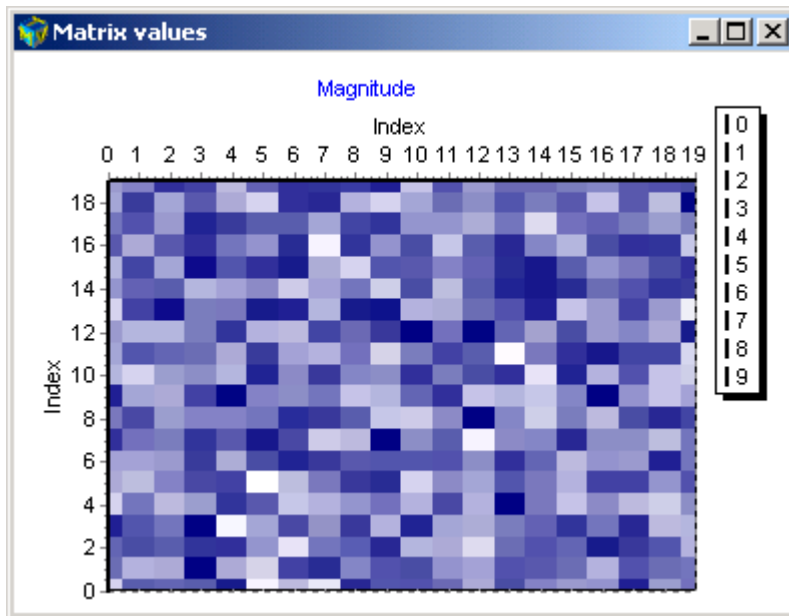


Figure 2

9 C++ Builder issues

9.1 Accessing array elements

MtxVec allocates memory for dynamic arrays in one continuous memory block. (fortran-style). This method is not compatible with two or more dimensional dynamic arrays allocated in C++. This type of memory allocation is necessary to make the arrays compatible with fortran libraries. (And use LAPACK). Because most of the matrix operations are based on LAPACK, this type of memory allocation is mandatory.

Consequently, in C++Builder matrix elements cannot access the elements via Values property (example: `a->Values[i][j] = 2;`, will not work), because sometimes this will generate range check errors. This range check errors are false alarms raised due to the memory allocation incompatibility issue.

In Delphi, T2DSampleArray type (used by TMtx) can be used also for direct pointer reference, because only the range checking mechanism is broken and, if range checking is turned off range, the app will run just fine. But this is not the case for C++Builder.

The C++ template for Delphi dynamic arrays used in CBuilder uses its own explicit range checking, which cannot be disabled and the T2DSampleArray type therefore cannot be used in C++ Builder to access the memory of TMtx.

To circumvent this problem, elements have to be accessed via a secondary "values" property:

```
a->values[i][j] = 2;
```

or

```
a->cvalues[i][j] = StrToCplx("2+i");
```

Delphi dynamic arrays are wrapped with a template in C++ Builder making each access very slow. By obtaining a pointer of P2DSArray type this can be avoided and full speed of can be obtained. (but no range checking) The P2DSArray type pointer can be conveniently obtained via the Enlist method.

With C++Builder it is recommended to use the Enlist method to access the elements of TMatrix or of TVector object. The Enlist method in C++Builder does not require a dismiss pair, as it is the case in Delphi.

The `a->values[i][j]` access is slower than a memory pointer, but still faster than the template used by CBuilder for Delphi dynamic arrays.

9.2 Open array parameters and SetIt method

One other important issue with C++Builder are open array parameters. Delphi allows passing arrays to routines without the need to first allocate them explicitly. Because many examples in MtxVec documentation use the SetIt routine an example below shows how the SetIt routine can be used in C++Builder:

```
//-----  
  
#include <vcl.h>  
#pragma hdrstop  
  
#include "Unit1.h"  
#include "MtxVecCPP.h"
```

```
#include "MtxVec.hpp"
#include "MtxVecTee.hpp"
#include "MtxVecEdit.hpp"
//-----
#pragma package(smart_init)
#pragma resource "*.dfm"
TForm1 *Form1;

//-----

void __fastcall TForm1::Button1Click(TObject *Sender)
{
    TVector *x, *y, *spec;
    TMatrix *a, *b, *c;

    CreateIt(y,x, spec);
    CreateIt(a,b,c);

    try
    {
//        Setting a matrix to a vector

        double M[4][4] = { {1 , -3 , 5, -3},
                           {-1 , 12 , 0.3, 2.5},
                           {5 , 1.22, 2.33, -0.5},
                           {2.4, -1 , 5, 3} };
        int FinalIndex = 4*4-1; //array is zero based.
        y->SetIt(false,&M[0][0],FinalIndex);

        ViewValues(y,"Vector",true);

//        Setting a matrix to a matrix

        a->SetIt(4,4,false,&M[0][0],FinalIndex);
        ViewValues(a,"Matrix",true);

//        Setting a vector to a matrix

        double Mv[16] = { 1 , -3 , 5, -3,
                          -1 , 12 , 0.3, 2.5,
                          5 , 1.22, 2.33, -0.5,
                          2.4, -1 , 5, 3};

        a->SetIt(4,4,false,&Mv[0],FinalIndex);
        ViewValues(a,"Matrix from Vector",true);

//        Getting values from a vector (pointer access)

        PArray darray; // equals to {double *darray;} if TSample is double
        Enlist(y,darray);
        darray[0] = darray[1];
        ViewValues(y,"First element changed",true);

//        Getting values from a matrix (pointer access)

        P2DSArray d2array; //equals to {double **d2array;} if TSample is
double
        Enlist(a,d2array);
```

```

        d2array[0][1] = d2array[1][2];
        ViewValues(a,"Second element changed",true);

        DrawValues(y,Series1);
    }
    _finally
    {
        FreeIt(x,y,spec);
        FreeIt(a,b,c);
    }
}
//-----

```

9.3 Default parameter values

In C++ Builder versions 4 and 5 the compiler does not correctly assign default values to procedure parameters. A routine declared like this

```
function CplxToStr(const Z: TCplx;const Digits: integer = 0): string;
```

can be called in Delphi like this:

```

var a: string;
      z: TCplx;
begin
    z := Cplx(2,3); // z = 2+3i
    a := CplxToStr(z); //a = "2+3i", Digits automatically become 0
end;

```

In C++Builder versions 4 and 5 all parameter values have to be typed out.

Delphi:

```

var a: string;
      z: TCplx;
begin
    z := Cplx(2,3); // z = 2+3i
    a := CplxToStr(z,0); //a = "2+3i"
end;

```

C++Builder:

```

{
    TCplx z;
    AnsiString a;
    z = Cplx(2,3); // z = 2+3i
    a = CplxToStr(z,0); //a = "2+3i"
}

```

While C++Builder 6 does correctly assign all default values to parameters, it does not display in the *code insight* that the routine does have default values assigned to some parameters. One other potential bug are also enumerated types. If variables of enumerated types are passed to the routine, the programmer must ensure that the variables are initialized. Out-of-range errors for enumerated types will not be caught in C++Builder.

10 Getting ready to deploy

Once the app has been debugged and is ready to be deployed, three files required by MtxVec have to be included in the distribution package. This files are located in the windows\system or winnt\system32 directory: mkl_support.dll, lapack.dll and bdspp.dll.

Three files form a minimum required dll set for MtxVec to run. If there is a need to speed things up by supporting specific CPU's (especially P4), there is a set of libraries located on www.dewresearch.com

If distribution size is a problem, we can make a build of custom size dll's for your specific application. Lapack dll's are linked dynamically and any routines not needed can be simply left out.

Since Intel MKL version 6.0, all CPU's are supported with a single dll: lapack.dll. mkl_support.dll is no longer required.

11 Debugging MtxVec

Delphi offers an excellent debugger. However there are a few things that have to be addressed separately.

11.1 Memory leaks

The programmer should make sure to always match the Create and Destroy methods of TObjects and CreateIt/Freelit pairs for TVec and TMtx objects, as already emphasized. MtxVec unit holds two global variables: MtxCacheUsed and VecCacheUsed. Their value will show the number of unfreed objects. After the application has finished using MtxVec routines, these two variables should have a value of 0. This means that all objects for which CreateIt was called, were also passed to the Freelit routine. For TObject type no such reference counting exists.

11.2 Memory overwrites

When using TVec and TMtx the memory overwrites can be difficult to notice. The reason for this is that TVec and TMtx objects residing in the object cache have preallocated a specified number of elements. Such pre-allocation speeds memory allocation for small vectors and matrices considerably and also makes reallocations faster. MtxVec explicitly checks all parameters passed to TVec and TMtx routines for range check errors. This however does not help, if the algorithm accesses the values directly from TVec.Values/TMtx.Values array. For TVec it helps, if range checking offered by the Delphi compiler is turned on. (Project - > Options - > Compiler - > Run time errors - > Range checking). When working with matrices, the range checking mechanism offered by the compiler raises false alarms unless TMtx.Values1D array is used. This false alarm occur, because the memory for matrices is allocated in a single block, which is not in compliance with the Delphi dynamic arrays. The memory for matrices must be allocated in a single block so that it is compatible with Fortran style memory allocation and LAPACK. Anyhow, to protect the program against memory overwrites when working with matrices, TMtx.Values1D array can be used for which the range checking mechanism works OK and the memory preallocation feature should be temporarily disabled. The memory preallocation is disabled by calling:

```
SetVecCacheSize(0, 0);  
SetMtxCacheSize(0, 0);
```

The first parameter defines the number of TVec/TMtx objects created in advance and the second parameter defines the number of array elements for which to preallocate the memory. By setting memory preallocation to zero AV's will be raised much closer to the actual cause of the problem.

11.3 Intel SPL error checking

MtxVec makes heavy use of Intel Signal processing library (Intel SPL) to support the latest CPU architectures (P4, Itanium etc...). The error checking mechanism of Intel SPL does not allow "silent" error reporting without some kind of memory corruption. This means that once an invalid parameter is encountered, the application calling the dll will be halted without any error message. Because MtxVec does its own checking of valid parameter combinations, this problem is minimized. In some cases however the user might pass an invalid parameter combination, which was not predicted. In such cases the application will simply terminate. You can use step by step debugging to see which routine caused the problem and prevent invalid parameters being passed to the dll routine.

MtxVec v1.5	1
<i>A programming guide to MtxVec</i>	1
1 Introduction	2
1.1 Indexes, ranges and subranges	2
1.2 TVec and TMtx methods as functions	3
1.3 Create and destroy	3
1.4 Complex data	4
1.5 MtxVec types	4
1.5.1 TSample.....	4
1.5.2 TCplx.....	4
1.5.3 TSampleArray, TCplxArray.....	5
2 Programming style	5
2.1 Try-finally blocks	5
2.2 Raising exceptions	6
2.2.1 Invalid parameter passed:.....	6
2.2.2 Reformat the exception.....	6
2.3 By all means indent code	6
2.4 Do not create objects within procedures and return them as result:	7
2.5 Use CreateIt/Freeit only for dynamically allocated objects whose lifetime is limited only to the procedure being executed.	7
3 Accessing values of TVec and TMtx	7
3.1 Default array property access	7
3.2 Dynamic array pointer	8
3.3 Direct dynamic array pointer	8
4 Memory management	9
5 Range checking	10
6 Why and how NAN and INF	10
6.1 Delphi 4 and Delphi 5 issues with NAN and INF	11
7 Serializing and streaming	12
7.1 Streaming with TMtxComponent	12
7.2 Streaming of TVec and TMtx	12
7.3 Write TMtx to a text file	12
7.4 Read TMtx from a text file	13
8 Input-output interface	13
8.1 Reading and writing raw data	13
8.2 Formatting floating point values – FloatToString	13
8.3 Printing current values of variables	14
8.4 Displaying the contents of the TVec and TMtx	14

8.4.1	As a delimited text.....	14
8.4.2	Within a grid.....	15
8.5	Charting and drawing.....	15
9	<i>C++ Builder issues.....</i>	17
9.1	Accessing array elements	17
9.2	Open array parameters and SetIt method	17
9.3	Default parameter values.....	19
10	<i>Getting ready to deploy.....</i>	20
11	<i>Debugging MtxVec.....</i>	20
11.1	Memory leaks.....	20
11.2	Memory overwrites.....	20
11.3	Intel SPL error checking.....	21