
Subject: Re: Object overhead

Posted by [John-David T. Smith](#) on Mon, 24 Sep 2001 17:26:05 GMT

[View Forum Message](#) <> [Reply to Message](#)

"Pavel A. Romashkin" wrote:

>
> I noticed that the size of an object array, where each object only has
> one empty pointer field, is 4 times larger than the size of an empty
> pointer array of the same length. PTRARR(1000), when saved, takes about
> 44 Kb. OBJARR(1000) of {junk, data:PTR_NEW()} takes 176 Kb. Why is that?
> Also, it takes twice the time (0.79 vs 0.44 s) to save an object array
> than the pointer array. It can not have to do with the volume of saved
> data because it takes 0.05 s to save an 80 Kb FLTARR(20000).
> I thought of converting a data processing program to objects but found
> that storage space required almost doubles (well, partially due to some
> redundancy in object fields) and save/restore time increased tenfold.
> For storing many thousands of data objects, this matters.
> Thank you,
> Pavel

Why don't you compare:

ptrarr(1000) to {junk_struct,data:ptrarr(1000)} to obj_new('junk') with

```
pro junk__define
  st={JUNK,data:ptrarr(1000)}
}
```

I'd suspect they'll be close, especially the last two. The basic issue is that arrays *of* structures, and arrays *in* structures must necessarily receive quite separate treatment, because structures are themselves such complex beasts.

Also, as Craig pointed out, this is an unfair test in itself, since the overhead of pointer arrays will be seen only when you actually attach them to some data, and will then likely dominate performance and size. As an example of this for the size issue, I tried:

```
IDL> a=ptrarr(1000) & for i=0,999 do *a[i]=findgen(3*i+1)
```

saving "a" yields 6107296 bytes on disk.

```
IDL> a=replicate({JUNK,data:ptr_new()},1000)
IDL> for i=0,999 do a[i].data=ptr_new(findgen(3*i+1))
```

which is 6107340 bytes on disk. Almost exactly the same! Then I tried:

```
function KNUJ::Init, data
```

```

    self.data=ptr_new(data)
    return,1
end

```

```

pro KNUJ__define
    struct={KNUJ,data:ptr_new()}
end

```

```
IDL> a=objarr(1000) & for i=0,999 do a[i]=obj_new('KNUJ',findgen(3*i+1))
```

which is 6239328 on disk.

I.e., around a 2% savings for the first two methods. Not terribly meaningful differences.

The only conclusion you can draw is that pointer array members carry about 45 bytes of meta-data each, and objects (and probably structures) carry about 180 bytes of meta-data each.

As an example, take structures. All C types of IDL variables are documented in external/export.h. A bit of piecing together reveals an array of structures looks like:

```

typedef struct {          /* IDL_VARIABLE definition */
    UCHAR type;           /* Type byte */
    UCHAR flags;          /* Flags byte */
    IDL_ALLTYPES value;
} IDL_VARIABLE;

```

A generic variable, not carrying too much baggage, about 10 bytes.

An array looks like:

```

typedef struct {          /* Its important that this block
                           be an integer number of longwords
                           in length to ensure that array
                           data is longword aligned. */
    IDL_MEMINT elt_len;   /* Length of element in char units */
    IDL_MEMINT arr_len;   /* Length of entire array (char) */
    IDL_MEMINT n_elts;    /* total # of elements */
    UCHAR *data;          /* ^ to beginning of array data */
    UCHAR n_dim;          /* # of dimensions used by array */
    UCHAR flags;          /* Array block flags */
    short file_unit;      /* # of assoc file if file var */
    IDL_ARRAY_DIM dim;    /* dimensions */
    IDL_ARRAY_FREE_CB free_cb; /* Free callback */
    IDL_FILEINT offset;   /* Offset to base of data for file var*/
}

```

```
IDL_MEMINT data_guard;    /* Guard longword */
} IDL_ARRAY;
```

Each structure itself looks like:

```
typedef struct {           /* Reference to a structure */
    IDL_ARRAY *arr;        /* ^ to array block containing data */
    struct _idl_structure *sdef; /* ^ to structure definition */
} IDL_SREF;
```

which is an array and the definition (tag names, etc.) The sdef field is "a pointer to an opaque IDL structure definition", i.e. they ain't gonna tell you.

What is the native size of a mostly empty IDL structure? Even without knowing the size of _idl_structure, we already have the following sizes (in bytes)

```
IDL_VARIABLE:  20
IDL_ARRAY:     64
IDL_SREF:      8
IDL_StructDefPtr: 4
+++++
          96 bytes + (_idl_structure)
```

Each tag name is associated with a struct of type IDL_STRUCT_TAG_DEF (16 bytes), so we're up to 112 bytes for at least one tag, not to mention the space for the tag name itself. Add in a few more pieces of meta-data inside of _idl_structure, and all of a sudden 180 bytes with 1 tag and without any data at all isn't inconceivable. So the moral of the story: IDL structures **weigh** much more than IDL scalars.

If you have ~100 bytes of data or so per entity, and can handle the loss of flexibility, you'll realize real space savings avoiding all the extra cruft attached to structures. If you have much more data per storage entity (struct, pointer array member, etc.), as in the above example, the space savings will be marginal.

JD
