

Safe arrays and pointers for C through compatible additions to the language

John Nagle

Discussion draft, round 3 – September, 2012

Introduction

Buffer overflows continue to plague C and C++ programs. This is a draft proposal for dealing with that problem. The basis of this proposal is to define means for expressing the size of arrays in C. C already has fixed-size arrays with useful semantics. In this proposal, the existing syntax for fixed-length arrays is generalized to allow known-length arrays, where the length of the array is determined at variable initialization time. With relatively minor and compatible changes to C, the most troublesome causes of program crashes and security vulnerabilities can be dealt with.

In any useful program, each array has a size known to the programmer. In C, there is currently no way to consistently express that size in the language. This proposal adds that capability.

A viable extension to C in this area must be compatible with existing code. This approach offers an optional “strict mode”. Translation units can be “strict” or “non-strict”. Existing code will compile in non-strict mode. Strict code can call non-strict code, and non-strict code can call strict code. Existing function declarations, such as widely used library interfaces, can be rewritten in a style which prevents array bounds overflows when called from strict code. Existing non-strict code can call these same interfaces. A gradual migration to strict code free of buffer overflows is thus supported.

Scope

An ISO document on common programming language vulnerabilities, “Information Technology — Programming Languages — Guidance to Avoiding Vulnerabilities in Programming Languages through Language Selection and Use”, ISO/IEC TR 24772, provides a standard taxonomy of flaws in programming language designs. The specific flaws in C addressed by this proposal are:

- 6.8 String Termination
- 6.9 Buffer Boundary Violation (Buffer Overflow)
- 6.10 Unchecked Array Indexing
- 6.11 Unchecked Array Copying
- 6.12 Pointer Casting and Pointer Type Changes
- 6.13 Pointer Arithmetic
- 6.14 Null Pointer Dereference
- 6.40 Type-breaking Reinterpretation of Data

These are the errors most closely associated with program crashes and low-level security vulnerabilities.

Not addressed are:

- 6.15 Dangling Reference to Heap
- 6.41 Memory Leak

Preventing errors of those types requires automatic memory management or extensive static analysis. Those flaws are difficult to fix without unacceptably large changes to the C language. Informally, a problem with adding garbage collection to C is that C is the language in which one writes the garbage collector.

This document uses as its working definition of the C language the Committee Draft of April 12, 2011, N1570, ISO/IEC 9899:201x.

Summary of language changes

The language changes required are summarized below.

Features imported from C++11

- References
- **auto**
- **decltype**

Relaxation of existing restrictions

- Expressions allowed in most array size declarations. (This is the key enhancement.)

New reserved words

- **lengthof** (*operator, similar to sizeof*)
- **force_cast** (*storage class, for casts only, allows forcing certain casts in strict mode*)
- **void_space** (*type, for arrays of memory with unspecified contents*)

New pragma

- **#pragma strict**

New standard macros

- *arraytype* **&(array_slice(array, start, end))[end - start];** (*Generic function as macro*)
- *typeofvalue* **init_space(variable, value)** (*Generic function as macro*)

New restrictions (non-strict mode)

- None on existing code other than avoidance of the new reserved words.

New restrictions (strict mode)

- Conversion from a pointer to an array to a reference to an array of known length requires an explicit **force_cast**, to indicate that the programmer has commanded an unsafe operation.
- Conversion from anything which can be NULL to a reference implies a run-time constraint check.
- Unions may not contain pointer or reference types, and all parts of a union must have the same length.
- To be modified by pointer arithmetic, a pointer variable must be initialized at declaration to an element of an array. Subscripting a pointer is allowed, but the result is read-only (“const”).
- The only allowed values for the pointer variable used in pointer arithmetic are ones which point to an element of the array to which it was originally assigned. (Pointing to an element one past the end is allowed, for backwards compatibility.)

Other

- The type of a string constant is a reference to an array, rather than a pointer to an array.

Rationale

Goals

- No change to run-time representation of data. (no descriptors or “fat pointers”)
- Minimal modifications to the C language.
- Extensions to be potentially compatible with C++11.
- Define “strict” and “non-strict” forms of C translation units; existing C99 code can be compiled as “non-strict” translation units and “strict” translation units can interoperate with non-strict translation units.
- If all translation units are of “strict” form, and no unsafe overrides are used, programs should be protected against the flaws listed above.
- Hidden memory management mechanisms (garbage collection, reference counting) should not be required.

Previous work

- Microsoft source code annotation language (SAL)

- Cyclone (AT&T research language)
- SCC, the Safe C compiler.
- Ccured
- MemSafe
- Incompatible dialects and related languages

Fixed-length arrays in C and C++

The basis of this proposal is to define means for expressing the size of arrays in C at all places where arrays are used. Most arrays in C and C++ are passed around as pointers, with no length information. There is, however, an exception – fixed sized arrays. C and C++ support fixed-length arrays. Fixed-length arrays are typed objects of known length. They can be defined as types, and, in C++, references to fixed length arrays are fully supported.

```
int arr100[100];    // an array of 100 ints
arr100& arr100r;  // ref to array of 100 ints
```

This is valid C++ today. C++ already supports the concept of a reference to an array. (Arrays of references are prohibited by the C++ standard, but references to arrays are not.)

Where the compiler knows the length of an array, we can access that information via **sizeof**. This is convenient, but only available for arrays of fixed length.

C99 variable length arrays

C99 added support for dynamically sized temporary arrays. This feature is often used in numerical code, where it is common to call a function which needs a temporary array for the duration of the function call.

```
float mathfn(const float in[ ], size_t n)
{
    float workarr[n];    // work array of length n
    ...
    return(result);
}
```

This was the first place in C where an array could be resized at compile time with the compiler aware of the size.

C99 also added variable length array parameters.

```
float fn(size_t n, double[n][n])
{
    ...
}
```

Variable length array parameters have pointer, rather than array, semantics within the function body. Per N1570 §6.7.6.3p7: *A declaration of a parameter as "array of `_type_`" shall be adjusted to "qualified pointer to `_type_`", where the type qualifiers (if any) are those specified within the [and] of*

the array type derivation. The effect of this conversion is that the first subscript of a variable-length array parameter is lost to the body of the called function. It is not accessible via **sizeof**. Contrast this with a local variable-length array, where **sizeof** returns the size occupied by the array. These semantics follow classic C semantics for array parameters of constant size or of unknown size.

Both of these features were mandatory in the C99 standard. However, they were not widely implemented in commercial compilers, and were little used in production code. (We have been unable to find a single instance of a variable-length array parameter in published open-source code. Five instances of local variable-length arrays in open-source code have been found to date, including one which was briefly in the Linux kernel but was removed.)

Reflecting this situation, N1570, at §6.7.6.2p4, reads: *"Variable length arrays are a conditional feature that implementations need not support"*.

Because of this history, the proposal here is to phase out C99-style variable length array parameters and replace them with a similar feature expressed as C++ style references. Variable-length local arrays can be retained as a feature.

Known-length arrays

The term “known-length array” here refers to a new approach to variable-length arrays for which all dimensions are known. The existing syntax and semantics of fixed-length arrays, C99 dynamically sized arrays, and C++ references can, with slight modifications, allow the consistent support of known length arrays. The syntactical change required is to allow expressions in a few places where, at present, only constants are allowed. It is proposed to allow expressions in array dimensions in type declarations in the following places:

- local variable and typedef declarations
- function parameter declarations
- casts within initialization expressions
- structure definitions (under certain restrictions)

When a known-length array is defined or passed as a parameter, its length is the value of the dimension expression evaluated at initialization time. This follows the usage for C99 variable length array parameters.

This feature allows a function such as the following:

```
void copybyref(size_t n, int (&a)[n], const int (&b)[n])
{
    for (int i = 0; i < lengthof(a); i++)
        { a[i] = b[i]; }
}
```

This is a size-safe copy, intended to be used as follows:

```
int a0[100], a1[100]; // defined with a type declaration
copybyref(100, a1, a0); // copy array of known size
```

```
int b0[50], b1[50];      // smaller arrays
copybyref(50, b0, b1);  // copy array of a different size
copybyref(100, b0, a0); // size mismatch - compile time error
```

The last “copybyref” call is an error, because the array sizes do not match. For a constant “n” this can be caught at compile time. For a variable n, it must to be a run time check. Handling of run time checks is covered below in the section on subscript and call checking.

If subscript and size constraint checking is enabled, the compiler must check, at a call to “copybyref”:

```
lengthof(a) == n
lengthof(b) == n
```

If the programmer calls “copybyref” as follows:

```
copybyref(lengthof b0, b0, a0);
```

The checks become

```
lengthof(b0) == lengthof(b0)
lengthof(a0) == lengthof(b0)
```

The first check is optimized out as an identity. The second check still requires a run-time check. But only one check per function call is required. This is far cheaper than general subscript checking.

Here, a run-time check is required only because two arrays must match. In many common cases, especially some of the classic causes of buffer overflows, there is no additional overhead. For example, the classic UNIX *read* call would now be expressed as:

```
int read(size_t n; int fd, void_space(&buf)[n], size_t n);
```

(The type `void_space` is discussed below; it's simply a type with a size of 1, usable for arrays, for which the bytes have no predetermined meaning. Think of it as `void*` with length. The initial “size_t n;” is a forward declaration, discussed below.)

The required constraint check to be generated by the compiler is

```
lengthof(buf) == n
```

If called with

```
int buf[512];
...
int stat = read(fd, buf, lengthof(buf));
```

the check required is

```
lengthof(buf) == lengthof(buf)
```

which implementations should optimize out as an identity.

Worth noting is that the function declaration of `read` above is, in non-strict mode, compatible with existing code which passes `buf` as a pointer to a char.

Why references?

A key part of this proposal is the use of C++ style references in C. As shown above, references to arrays can contain size information about the array. Pointers to arrays in C refer to the type of the first element, with no array length information. Changing the semantics of pointers in C would break existing code. Adding references to the C language will not break anything; C++ already has both references and pointers. The relationship between the two is well understood and not a cause of trouble, so this can be done with confidence.

C99 already allows variable declarations to be interspersed with executable code. References must be initialized in C++, and this restriction would be maintained with references in C. This proposal allows expressions in array sizes in declarations, and those expressions must be evaluated where the declaration appears.

```
int b[100];           // array of 100 ints
...
int (&rb)[lengthof(b)] = b; // ref to array b
```

Assignments to references are subject to type checking. Where the reference on the left hand side has a known length determined at run time, this may imply a run-time constraint check on array size. This is considered a subscript check, as defined below.

Strict mode

The enhancements to C syntax and semantics described here are backwards compatible with existing C code. Beyond this is “strict mode” which disallows some unsafe constructs and operations. Strict mode requires known size arrays in situations where overflow is possible. “Strict mode” is defined on a per-source-file basis. Strict and non-strict code can interoperate.

New rules enforced in strict mode:

- Conversion from a pointer to an array to a reference to an array of known length requires an explicit **force_cast**, to indicate that the programmer has commanded an unsafe operation.
- Conversion from anything which can be NULL to a reference implies a run-time constraint check.
- Unions may not contain pointer or reference types, and all parts of a union must have the same length.
- To be modified by pointer arithmetic, a pointer variable must be initialized at declaration to an element of an array. Subscripting a pointer is allowed, (both with “[]” and with “p + offset” notation) but the result is read-only (“const”). This last is required for compatibility with existing null-terminated string usage.
- The only allowed values for the pointer variable used in pointer arithmetic are ones which point to an element of the array to which it was originally assigned. (Pointing to an element one past the end is allowed, for backwards compatibility.)

Programs written in strict mode should be protected against the list of memory safety errors covered in the “scope” section above. This should eliminate most buffer overflow problems.

Strict mode for a compilation unit is indicated by

```
_Pragma("strict")
```

or

```
#pragma strict
```

When non-strict code calls strict code, or vice versa, known-length array type checking is performed according to the following table. Here, T is a type, and n is a length parameter for a known-length array.

Declaration	Call	Types match in strict mode?	Types match in non-strict mode?	Notes
T var[n]	T var[n]	Yes	Yes	
T var[n]	T var[] T* var	No	Yes	
T* var	T* var	Yes	Yes	In strict mode, var cannot be subscripted in the function body.
T var[]	T var[]	No	Yes	“var[]” empty length form not allowed in strict mode.

The intent of these rules is to allow strict-mode declarations such as

```
int read(size_t n; int fd, void_space[n], size_t n);
```

to be called from non-strict code with an common old-style call such as

```
char buf[100];  
int fd;  
...  
int stat = read(fd, buf, 100);
```

The array dimensions are not checked in this backwards-compatible case. The intent is to allow converting libraries, especially standard APIs, to the strict form, while allowing both strict and non-strict code to call them.

This feature does not extend to allowing non-strict code to pass an array of pointers to a multidimensional array of known length. So the classic declaration form

```
void gaussj(float **a, int n, float **b, int m);
```

is not interchangeable with the strict form

```
void gaussj(int m; float (a&)[n][n], int n, float (&b)[n][m],  
int m);
```

because **a** and **b** are arrays of pointers.

Implications of known-length arrays

With known-length arrays and strict mode introduced, their use and consequences can be examined.

How array size information is stored

No array descriptors are generated by the compiler. The programmer tells the compiler the size of the array as an expression, and that expression is evaluated at the point in the program where the relevant declaration appears.

```
int asize = 100;    // size wanted
int atab[asize];   // declare known-length array
asize = 200;       // this does not resize the array
printf("Length of atab is %d\n", lengthof(atab)); // prints 100
```

In the above worst case example, the compiler must generate a temporary is needed to store the size of the array, because the inputs to the expression defining its size change. This is usually unnecessary.

A struct with an array as the last component

An especially useful idiom is a structure with an array as the last component. C supports this now, but the size of the array is always undefined, expressed with []. This proposal allows such arrays to be of known size. The array, and the structure, then have a known size. and

A simple example is

```
typedef struct msgitem {
    const size_t len;
    char itemvalue[len];
};

struct msgitem firstmsg = { 100, 0 }; // empty msgitem, size 100
```

The length, **len**, is an element in the same structure. This ties the length of the array to a value which can be found from the structure. This tie is a property of the type declaration. The general rule for such declarations is that the dimension expression is evaluated in the context of the struct. Fields of the struct may appear in the expression. Such structs should be initialized as a unit, with an initialization expression.

For such types, the size of the type, and the length of the array, are known to the language. They can be accessed with **sizeof** and **lengthof**. If subscript checking is enabled, lengths are checked on assignments to and from the entire variable.

The examples in Appendix 2 show a safe string type implemented using this approach.

Associating length with an array pointer in a structure

Alternatively, a user might want a structure which carries both the array and its size. The pointer-based form is

```
typedef str { size_t len, char* data};
```

New reference form:

```

typedef str { size_t len, char (&data)[len]};
...
str s1 = {n, MALLOC(char, n) }; // structure initialization

```

The key point here is that `char (&data)[len]` is a reference to a string of known length. As in the previous example, `len`, is an element in the same structure, and ties the length of the array to a value which can be found from the structure. Multiple arrays may be present in the same structure. However, all such arrays must be allocated during structure initialization, because a reference cannot be null.

Array length expressions

C99 variable-length array parameters were allowed arbitrary expressions in their array length declarations. This created potential incompatibility between function prototypes and function definitions. The solution used in C99 was to permit differences between the two which did not affect the type of the pointer produced after the array passed had been reduced to a pointer type without length information.

This proposal is intended to insure that array length at a function call is consistent with array length when the function is entered. Thus, it must be insured that evaluation of an array length expression produces consistent results in any context where it can be evaluated.

An additional requirement, for future compatibility with C++, is that it must be possible to express all the type information associated with a function prototype as a unique string for the purpose of resolving function overloading. (This refers to “name mangling” in C++).

It is thus necessary to limit the form of array length expressions.

- In function prototypes, the only variables allowed in array length expressions are formal parameter variables declared within the function prototype. This is the first scope searched.
- In structure definitions, the only variables allowed in array length expressions are elements of the same structure definition. This is the first scope searched.
- Constants, including named constants, are allowed in array length expressions. Such constants must be evaluated at compile time.
- User-defined functions are not allowed in array length expressions.

These restrictions allow the common use cases, while permitting length checks between function call and function definition. Such checks are valid even across translation unit boundaries, which, of course, is a major point of this proposal.

It is sometimes necessary to use a formal parameter in a length expression which appears before the declaration of that parameter. A classic C convention, used both in the standard libraries and in many common operating system interfaces, is to declare functions with the length information after the size of a buffer or string. This is too important a convention to change. A classic example, from POSIX, is:

```
int read(int fd, char buf[n], size_t n);
```

The safe form of this, using the extensions in this proposal, is

```
int read(int n; int fd, char (&buf)[n], size_t n);
```

The first “size_t n;” is not a parameter; it is a declaration of n in a very local scope. This is a nonstandard extension from the Gnu Compiler Collection. GCC allows declarations in expressions generally, but for the purpose of this proposal, it is suggested that this feature be limited to forward declarations of formal parameters.

Note that the prototype for **read** above can be called either with a pointer (unsafely) or an array reference (safely). Existing code would be able to call a library with the new form of prototype. In “strict mode” translation units, only the second form, with size checks, would be allowed.

Type void_space

C uses the idiom **void *** for a pointer to data of undetermined type. Conversion to and from void without coercion is allowed. However, an array of **void** is not allowed; **void** has no size. A way is needed to talk about arrays of space without implying type. So it is proposed to add the built-in type **void_space**.

A variable of **void_space** is essentially the machine allocation unit of storage. This is a byte on almost all current hardware. (9 bit quarter words on Unisys ClearPath 36-bit mainframes.) The size of **void_space** is 1. An array of **void_space** is an array of bytes. Any variable, array, or struct that does not contain pointers or references can be converted to or from an array of **void_space**, without a cast, if the sizes match. A reference to an array of **void_space** can be converted to **void ***, and in non-strict code, the reverse is allowed.

Thus, any function which now uses a parameter of **void*** to represent an array can instead use a reference to an array of **void_space**. For example,

```
int read(int fd, void_space (&)buf[n], size_t n);
```

defines the standard POSIX/Unix read function in a way which is size-safe when called from strict code, and callable from non-strict code. This preserves programmer expectations about these common operations while enforcing size safety in strict code.

Type **void_space** values are converted to type **unsigned int** if used in an arithmetic context.

Casts

The syntax of casts follows that for declarations. As with reference declarations, the size must be evaluated at the point of the cast. Casts with expressions in array sizes are allowed in initialization expressions.

Ordinary casts must be memory-safe in strict mode. Unsafe casts must be expressed using some specific syntax, as with **reinterpret_cast** in C++. Unfortunately, the syntax of **reinterpret_cast** uses the corner-bracket template notation of C++, and is inappropriate for C. Thus, a syntax compatible with C is needed.

This, like most syntax issues, is a politically touchy problem. A solution should be

- Consistent with the expectation of C programmers.
- Backwards compatible with existing code
- Findable by simple searches to enable easy code auditing.

One option is a special storage class, used only in casts. A phrase such as **forced_cast** might be appropriate:

```
int val = 1234;
char* ptr = (forced_cast char*)val;    // UNSAFE
```

Even in strict mode, this would be required only when forcing a conversion which could impact memory safety. So such casts would be quite rare.

A few casts can generate run-time checks. Conversion of a pointer to a reference must generate an constraint check. Such checks are enabled when subscript checking is enabled.

Dynamic allocation (“malloc” and its friends)

This is one of the more difficult problems to handle safely. Yet, as it is frequently a source of program bugs of the buffer overflow variety, it is necessary to do so. The problems which must be addressed include:

- The classical return type of **malloc**, “void*”, cannot express size information.
- The return type from memory allocation must be convertible to other types.
- Because **malloc** can return NULL, it is not permissible to directly generate a reference from **malloc**.
- Pointers to references are not allowed.
- Structures which contain references must be initialized.
- The syntax for allocating memory must be concise.
- Exceptions are not available in C.
- Backwards compatibility with existing code must be maintained in non-strict mode.

All those constraints, taken together, create a tough problem. A few options, which should be viewed as illustrative, not part of a proposed library:

- Allocation functions which call some well known global function on failure. This impacts program structure, but provides the cleanest syntax.

```
void_space (&smalloc(size_t len))[len];
```

With this option there's is way to handle the error condition in in-line code. It's worth providing this for routine applications which do not need sophisticated out-of-memory handling.

- Allocation functions with a failure callback. This is for programs which need to explicitly handle out of memory conditions

```
void_space (&smalloc_failcheck(
    size_t len,
    void (*outofmemfn)(size_t)))[len];
```

The *outofmemfn* function, if non-null, would be called on an out of memory condition. If it is able to relieve the memory shortage and return, the program can continue. If not, the function should escape the call, using a **longjmp**, or terminate the program after any necessary closeout

actions.

- Return a pointer to a structure which contains a reference, or return NULL. The string type example in the appendix uses this approach.

Initialization of dynamically allocated space

In C, “malloc” normally provides memory space containing old semi-random data. For most data types, this is not a safety problem. However, for pointer safety, we must insure that allocated items with pointers or references are initialized. Zeroing new memory space is not an option for references, and because we're encouraging the use of references in place of pointers, this has to be done right.

So, when a new struct or array or structs is allocated, and the struct type contains pointers or references, it must be initialized from a valid initial value of that type. It would be preferable to do this without explicit casting.

As C lacks generic functions other than built-ins, a type-safe way to encapsulate this is needed. A suitable macro, **init_space**, is suggested. This is a size and type safe replacement for the classic “memcpy”.

```
typeofvalue init_space(variable, value)
```

This can be implemented easily with

```
#define init_space(var, val)\
(assert(sizeof(val) == sizeof(var)),\
memcpy(&(var),\
&(val), sizeof(val)),\
(decltype(val)&) (var))
```

This simply copies *value* into *variable*, with type and size checking. The return type is the type of *value*. (This is why we need **decltype**; to obtain the type for such casts.) The *variable* parameter can be an array of known or fixed size. This is needed because array assignment isn't allowed, and assignment to an array reference assigns the reference, not the content. In C++, a template function would be used, but in C, a macro is required.

Example usage, where a structure with some references is to be initialized, is

```
struct bufferpair {\
    size_t len;\
    char (&buf1)[len];\
    char (&buf2)[len];\
};\
...\
// allocbufferpair - allocate a buffer pair\
struct bufferpair* allocbufferpair(size_t len)\
{\
    const itemproto bp = {len, smalloc(len), smalloc(len)};\
    bufferpair& p = init_space(smallloc(sizeof(itemproto)),\
        itemproto);\
    return(bufferpair);\
}
```

```
}
```

A C initializer is used to construct the object, which is then copied to the newly allocated space. No casts are required, and this is both type and size safe. The **malloc** function is from the previous example.

Impact on other language features

The basic proposal has now been described. The impact of the proposal on various features of the C language is next.

Unions

In strict mode, elements of unions are restricted to non-pointer and non-reference types. Unlike struct types, union types may not have arrays of known but variable length.

In most existing code, unions are used primarily for data being passed through networks, files, and message passing systems. Such data, since it comes from external sources, seldom contains pointers. Adding language support for safe discriminated variant records, as was done in Cyclone, seems excessive for the infrequency with which it would be used.

In strict mode, the only way to override pointer type safety is through the explicit use of **force_cast**. This makes it easy to search and audit programs for pointer type problems.

Multidimensional arrays

Multidimensional arrays of known size are supported. This simply generalizes the existing rules for multidimensional fixed-size arrays.

```
size_t order = 3;
...
float tab[order][order]; // square array of order N.
```

As with arrays of fixed size, these are not arrays of pointers, but dense arrays of elements. This will be a great convenience in numerical work. There is also a performance gain. On modern CPUs, multiplying by row size to compute a row offset is faster than accessing a table of pointers.

Pointer arithmetic

Pointer arithmetic is a standard idiom of C programming. Although it would be safer to eliminate it from the language, that isn't feasible for an update to C.

The most common uses of pointer arithmetic use local variables which are bound to a single array throughout their life. Such pointers behave much like C++ iterators. So we can allow pointer arithmetic under much the same restrictions applied to iterators. These are:

- The pointer variable must be initialized at its declaration to an element of an array.
- The only allowed values for the pointer variable are ones which point to an element of the array to which it was originally assigned. (Pointing to an element one past the end is allowed, for backwards compatibility.)

- The array must not have a scope narrower than that of the pointer variable, so that the array cannot go out of scope while the pointer variable is still in scope.

These restrictions apply only to pointer variables to which pointer arithmetic operations are applied, and apply only in “strict mode”.

Pointers with such properties are safe for pointer arithmetic purposes, provided there is range checking. Range checking for such pointers is possible without “fat pointers”, because the array to which they are bound is known at compile time.

Unsafe pointer arithmetic is allowed only through **force_cast**.

This allows a strict mode version of the classic C idiom:

```
void bcopy(char (&dst)[len], char (&src)[len], size_t len)
{   while (len--) { *dst++ = *src++; } }
```

Called from strict mode, the lengths of the parameters must match **len**, and the function itself has enough information to allow run-time subscript checking. Called from non-strict mode, this is compatible with the classic **bcopy** function.

If this were written with a **for** loop, subscript checks could be hoisted out of the loop, recognized as an identity, and eliminated, resulting in safe code without run time checks within the loop. On some platforms, one **mov** instruction could perform the copy without loss of memory safety.

String constants

String constants will now have type **char (&)[length]** or **wchar_t(&)[length]**. These types will convert to **char*** or **wchar_t*** for compatibility with non-strict code. The **sizeof** and **lengthof** operators can be applied and count the usual trailing null.

Null-terminated strings are a problem. A workable solution is to allow read access to const arrays without subscript checking. This allows reading junk, but not overwriting other variables.

Standard library safety

Library functions can be categorized as follows:

- Can be made fully compatible and safe with arrays of known size – **read**, **write**, **fread**, **fwrite**, etc.
- Unsafe and already deprecated – **gets**, **sprintf**, etc. These would be unavailable in strict mode.
- Safe for write, but not for read – **printf**, **snprintf**, **sprintf_s**, etc. Permitted in strict mode on the grounds that the impact of disallowing them is too high.
- Unsafe – Unavailable in strict mode. However, unsafe functions in non-strict translation units can be called from strict mode code.

Library issues are beyond the scope of this document, but appear to be manageable.

Size and subscript checking

With enough information available to perform size and subscript checking, implementations may offer

run-time checking as an option. The question is what to do when a run time error is detected. N1570, Annex K, provides a standard action to be taken – a call to a runtime-constraint handler function set via **set_constraint_handler_s**.

Three types of run-time checks are required:

- Function call array size checks. These checks are made, when necessary, in the calling translation unit. The function prototype provides enough information to check array sizes.
- Subscript checks.
- Pointer-to-reference conversion null pointer checks.

It is strongly recommended that implementations which do size and subscript checking optimize it. As a minimum, the following is suggested:

- common subexpression optimization should be applied to subscript checks within loops.
- Within “for” loops, checks should be hoisted to the beginning of the loop and done once per loop if at all possible.
- Hoisted checks which resolve to constant expressions or identities should be evaluated at compile time, and reported as errors (if failing) or optimize out (if succeeding).
- For function call array checks, identities should be recognized and **lengthof** should be understood. In particular, the idiom

```
int read(size_t n; int fd, char (&buf)[n], size_t n);
...
char buf[1024];
int stat = read(infilefd, inbuf, lengthof inbuf);
```

should be recognized by compilers as an identity and optimized out, eliminating all extra function call overhead for safe mode.

The goal of the above is to optimize out most, if not all, subscript checks in inner loops of common matrix calculations, and to eliminate unnecessary checking overhead for function calls.

It is explicitly permitted to detect and act on a subscript error as soon as it becomes inevitable. In other words, if an implementation can detect that a subscript error will occur on some iteration of a loop, and there is no way to exit the loop before reaching that iteration, the subscript error can be reported at entry to the loop. A potential exit from the loop via “break” or “return” must inhibit this optimization. A potential exit via “longjmp” or “exit”, which might be hidden in a called function, would not be considered to inhibit this optimization. This allows hoisting subscript checks to the top of loops.

Convenience features

With the language aware of array length, some convenience features can be added to make programming easier. These are all optional from the programmer's perspective, and need not be retrofitted to existing code. None of these features are essential to the proposal.

“lengthof”

C has a built-in “**sizeof**” operator. It's common to define **lengthof**, returning the number of elements of an array, as a macro. It is proposed to make **lengthof** a standard. Both **sizeof** and **lengthof** are defined for arrays of known size. This makes the size of an array immediately and reliably available to code.

Adding a new, short reserved word is always a difficult issue. This proposal does not change the existing de-facto standard semantics of **lengthof** as used in library macros.

“auto” storage type as in C++11.

“auto”, from C++11, is added to C. The motivation for this is the long declarations required for references to arrays of known length. Writing strict mode code requires that every pointer and reference declaration have full size information. This increases code verbosity and programmer workload unless some assistance is provided.

Iteration over an array

Almost all languages defined in the last 20 years (with the notable exception of JavaScript) have syntax for iterating over all the elements of a collection. For arrays of known size, the compiler now has enough information to provide such a construct. The syntax of this may be controversial, and needs discussion. Some options include

```
for (auto x in arr) { ... }
for (auto x forall arr) { ... }
for (auto x: arr) { ... }
```

Making “in” a reserved word impacts existing code, but yields the nicest syntax. Using this feature with “auto” leads to concise loop statements which are not error prone. At this time, no specific proposal is made in this area. Suggestions are welcome.

Subarray access

At times, it is desirable to talk about a subarray of an array. The existing C syntax for this is

```
char arr[100]
char p[] = &arr[20];
```

This tends to be error-prone, as the size of the subarray is rather vague. So it is proposed to borrow the *slice* operation from Javascript:

```
char arr[100];
char (&p)[80] = array_slice(arr, 20, 100);
```

This can implemented as a macro. The result of **array_slice** is a reference to a subarray of the array.

Following the Javascript and Python convention, the limits of slice are the start index and the end index + 1. Negative, overlapping, and out of bounds indices are errors and should be detected if subscript checking is enabled.

Another option is to borrow Python's subarray syntax:

```
Char arr[100];  
char (&p)[80] = arr[20:100]; // reference to a subarray of arr
```

This is concise, but more controversial, as it involves new syntax. It has the advantage that subarray references could be allowed on the left side of an assignment.

Conclusion

This is a workable solution to a decades-old practical problem, a problem which has probably caused more computer program crashes than anything else in the history of computing.

Alternatives

Making C a memory safe language without unduly changing the language is difficult. This proposal is intended as a direction for standard C, and as such, must be quite conservative in the changes proposed.

Annotation systems, such as Microsoft's SAL, can be effective if used. Because they're not integrated into the language, they tend to make programs more verbose. Thus, they're rarely used without strong managerial pressure.

Cyclone has a reasonable set of extensions, but strays far enough from C that it is a different language. Cyclone offers three types of pointers - "normal", "never-null", and "fat". New symbols ("@" and "?") are used to designate the different types of pointers. "Fat" pointers imply hidden run-time machinery. This is a valid approach, but it is not C.

SCC, the Safe C compiler, is claimed to be able to detect all pointer and access errors at run time. However, it required using an elaborate data structure for every pointer and substantial run-time overhead.

C++ addresses safety by trying to hide the unsafe constructs of C under a layer of standard templates. In practice, the templated constructs tend to "leak" raw C pointers, and buffer overflows in C++ applications are still common.

Other variants such as Objective-C and Java are useful languages, but they are not C.

Limitations

This approach still doesn't solve the dangling-pointer problem. There's no good way to do that without more memory management machinery behind the scenes. The goal here was to avoid adding hidden machinery.

Appendix 1 – Examples

From the UNIX/POSIX API:

Original function declarations:

```
int read(int fd, char* buf, size_t n);
```

New form:

```
int read(size_t n; int fd, void* buf, size_t n);
```

This is size-safe when called from new code, and compatible with old code that passes a “char*” as a parameter.

(More to follow)

Appendix 2 – an illustrative string library

This is an illustration of how a safe string library, usable in strict mode, might be constructed. It is not part of the proposal to specify such a library at this time.

```
//  
// sstring -- strict string  
//  
// A length, and a buffer. The content is null-terminated and  
// the length is the buffer size.  
//  
struct sstring {  
    size_t len;  
    char s[len];  
};  
  
//  
// sstring_alloc -- Safe string from null-terminated string  
//  
struct sstring* sstring_alloc(size_t len)  
{  
    // Create prototype of dummy string  
    struct sstring proto = { len, {0}};  
    struct sstring* dest = // use prototype to calculate size
```

```

        (struct sstring*) malloc(sizeof(proto));
    if (dest == NULL) return(NULL);    // fails on out of memory
    *dest = proto;                    // int from prototype
    return(dest);                     // return pointer to string struct
}

//
// sstring_make -- Safe string from null-terminated string
//
struct sstring* sstring_make(const char* src)
{
    const size_t srcnt = strlen(src); // copy size, including null
    // Create prototype of dummy string
    struct sstring proto = { srcnt+1, {0}};
    struct sstring* dest = // use prototype to calculate size
        (struct sstring*) malloc(sizeof(proto));
    if (dest == NULL) return(NULL);    // fails on out of memory
    dest->len = proto.len;              // copy length
    // Copy contents, including trailing null
    for (size_t i=0; i <= dest->len; i++) dest->s[i] = src[i];
    return(dest);                      // return pointer to string struct
}

//
// sstring_cat - concatenate safe string
//
struct sstring (&sstring_cat(sstring &dest, const sstring &src)
{
    const size_t destcnt = strlen(dest.s); // content length in dest
    const size_t srcnt = strlen(src.s);   // content length in src
    assert(destcnt + srcnt < dest.len);    // will it fit?
    for (size_t i = 0; i <= srcnt; i++)    // copy contents w/null
        { dest.s[i+destcnt] = src.s[i]; }
    return(dest);                          // return pointer to string struct
}

//
// sstring_del - delete string
//
void sstring_delete(sstring* dest)
{
    if (!dest) return;                    // if already deleted, done
    free(dest);                           // free string
}

//
// demo - strict string test demo
//
void demo(int fd)
{
    sstring* s1 = sstring_make("Hello");

```

```
sstring* s2 = sstring_make(" ");
sstring* s3 = sstring_make("World");
sstring* wrk = sstring_alloc(25); // create a working string
sstring_cat(wrk,s1); // add some words
sstring_cat(wrk,s2);
sstring_cat(wrk,s3);
// Write the string out.
// Type checking insures the sizes match.
// "sizeof" is meaningful for this type.
int stat = write(fd, wrk, sizeof(*wrk));
// Print the string using printf.
// No size checking here, but this is a const usage.
printf("%s\n", wrk->s);
}
```