

Paradyn Parallel Performance Tools

DynC API Programmer's Guide

13.0 Release
February 2024

Computer Sciences Department
University of Wisconsin–Madison
Madison, WI 53706

Computer Science Department
University of Maryland
College Park, MD 20742

Email dyninst-api@cs.wisc.edu
Web <https://github.com/dyninst/dyninst>



Contents

| | | |
|----------|---|-----------|
| 1 | DynC API | 2 |
| 1.1 | Motivation | 2 |
| 1.1.1 | Dyninst API | 2 |
| 1.1.2 | DynC API | 4 |
| 1.2 | Calling DynC API | 4 |
| 1.3 | Creating Snippets Without Point Information | 5 |
| 2 | DynC Language Description | 5 |
| 2.1 | Domains | 5 |
| 2.2 | Control Flow | 6 |
| 2.2.1 | Comments | 6 |
| 2.2.2 | Conditionals | 6 |
| 2.2.3 | First-Only Code Block | 6 |
| 2.3 | Variables | 7 |
| 2.3.1 | Static Variables | 8 |
| 2.3.2 | An Explanation of the Internal Workings of DynC Variable Creation | 8 |
| 2.3.3 | Creating Global Variables That Work With DynC | 9 |
| 2.3.4 | Data Types | 10 |
| 2.3.5 | Pointers | 10 |
| 2.3.6 | Arrays | 10 |
| 2.4 | DynC Limitations | 10 |
| 2.4.1 | Loops | 10 |
| 2.4.2 | Enums, Unions, Structures | 10 |
| 2.4.3 | Preprocessing | 11 |
| 2.4.4 | Functions | 11 |
| A | The Dyninst Domain | 12 |

1 DynC API

1.1 Motivation

Dyninst is a powerful instrumentation tool, but specifying instrumentation code (known as an Abstract Syntax Tree) in the `BPatch_snippet` language can be cumbersome. DynC API answers these concerns, enabling a programmer to easily and quickly build `BPatch_snippets` using a simple C-like language. Other advantages to specifying `BPatch_snippets` using dynC include cleaner, more readable mutator code, automatic variable handling, and runtime-compiled snippets.

As a motivating example, the following implements a function tracer that notifies the user when entering and exiting functions, and keeps track of the number of times each function is called.

1.1.1 Dyninst API

When creating a function tracer using the Dyninst API, the programmer must perform many discrete lookups and create many `BPatch_snippet` objects, which are then combined and inserted into the mutatee.

Look up `printf`:

```
1  std::vector<BPatch_function *> *printf_func;
   appImage->findFunction("printf", printf_func);
   BPatch_function *BPF_printf = printf_func[0];
```

Create each `printf` pattern:

```
   BPatch_constExpr entryPattern("Entering %s, called %d times.\n");
2  BPatch_constExpr exitPattern("Exiting %s.\n");
```

For each function, do the following:

Create snippet vectors:

```
   std::vector<BPatch_snippet *> entrySnippetVect;
   std::vector<BPatch_snippet *> exitSnippetVect;
```

Create the `intCounter` global variable:

```
   appProc->malloc(appImage->findType("int"), std::string("intCounter"));
```

Get the name of the function:

```
   char fName[128];
   BPatch_constExpr funcName(functions[i]->getName(fName, 128));
```

Build the entry `printf`:

```
   std::vector<BPatch_snippet *> entryArgs;
   entryArgs.push_back(entryPattern);
3  entryArgs.push_back(funcName);
   entryArgs.push_back(intCounter);
```

Build the exit printf:

```
1      std::vector<BPatch_snippet *> exitArgs;
      exitArgs.push_back(exitPattern);
      exitArgs.push_back(funcName);
```

Add printf to the snippet:

```
2      entrySnippetVect.push_back(BPatch_functionCallExpr(*printf_func, entryArgs));
      exitSnippetVect.push_back(BPatch_functionCallExpr(*printf_func, exitArgs));
```

Increment the counter:

```
      BPatch_arithExpr addOne(BPatch_assign, *intCounter,
                              BPatch_arithExpr(BPatch_plus, *intCounter, BPatch_constExpr(1)));
```

Add increment to the entry snippet:

```
      entrySnippetVect.push_back(&addOne);
```

Insert the snippets:

```
      appProc->insertSnippet(*entrySnippetVect, functions[i]->findPoint(BPatch_entry));
      appProc->insertSnippet(*exitSnippetVect, functions[i]->findPoint(BPatch_exit));
```

1.1.2 DynC API

A function tracer is much easier to build in DynC API, especially if reading dynC code from file. Storing dynC code in external files not only cleans up mutator code, but also allows the programmer to modify snippets without recompiling.

In this example, the files `myEntryDynC.txt` and `myExitDynC.txt` contain dynC code:

```
// myEntryDynC.txt
static int intCounter;
3 printf("Entering %s, called %d times.\n", dyninst'function_name, intCounter++);

// myExitDynC.txt
2 printf("Leaving %s.\n", dyninst'function_name);
```

The code to read, build, and insert the snippets would look something like the following:

First open files:

```
FILE *entryFile = fopen("myEntryDynC.txt", "r");
FILE *exitFile = fopen("myExitDynC.txt", "r");
```

Next call DynC API with each function's entry and exit points:

```
BPatch_snippet *entrySnippet =
    dynC_API::createSnippet(entryFile, entryPoint, "entrySnippet");
3 BPatch_snippet *exitSnippet =
    dynC_API::createSnippet(exitFile, exitPoint, "exitSnippet");
```

Finally insert the snippets at each function's entry and exit points:

```
1 appProc->insertSnippet(*entrySnippet, entryPoint);
  appProc->insertSnippet(*exitSnippet, exitPoint);
```

1.2 Calling DynC API

All DynC functions reside in the `dynC_API` namespace. The primary DynC API function is:

```
BPatch_Snippet *createSnippet(<dynC code>, <location>, char * name);
```

where `<dynC code>` can be either a constant c-style string or a file descriptor and `<location>` can take the form of a `BPatch_point` or a `BPatch_addressSpace`. There is also an optional parameter to name a snippet. A snippet name makes code and error reporting much easier to read, and allows for the grouping of snippets (see section 2.3.2). If a snippet name is not specified, the default name `Snippet_[<#>]` is used.

The location parameter is the point or address space in which the snippet will be inserted. Inserting a snippet created for one location into another can cause undefined behavior.

| <dynC code> | Description |
|------------------------------|--|
| <code>std::string str</code> | A C++ string containing dynC code. |
| <code>const char *s</code> | A null terminated string containing dynC code |
| <code>FILE *f</code> | A standard C file descriptor. Facilitates reading dynC code from file. |

Table 1: `createSnippet(...)` input options: dynC code

| <location> | Description |
|--|--|
| <code>BPatch_point &point</code> | Creates a snippet specific to a single point. |
| <code>BPatch_addressSpace &addSpace</code> | Creates a more flexible snippet specific to an address space. See Section 1.3. |

Table 2: `createSnippet(...)` input options: location

1.3 Creating Snippets Without Point Information

Creating a snippet without point information (i.e., calling `createSnippet(...)` with a `BPatch_addressSpace`) results in a far more flexible snippet that may be inserted at any point in the specified address space. There are, however, a few restrictions on the types of operations that may be performed by a flexible snippet. No local variables may be accessed, including parameters and return values. Mutatee variables must be accessed through the `global` domain.

2 DynC Language Description

The DynC language is a subset of C with a `domain` specification for selecting the location of a resource.

2.1 Domains

Domains are special keywords that allow the programmer to precisely indicate which resource to use. DynC domains follow the form of `<domain>`<identifier>`, with a back-tick separating the domain and the identifier. The DynC domains are as follows:

| Domain | Description |
|----------------------|---|
| <code>inf</code> | The inferior process (the program being instrumented). Allows access to functions of the mutatee and it's loaded libraries. |
| <code>dyninst</code> | Dyninst utility functions. Allows access to context information as well as the <code>break()</code> function. See Appendix A. |
| <code>local</code> | A mutatee variable local to function in which the snippet is inserted. |
| <code>global</code> | A global mutatee variable. |
| <code>param</code> | A parameter of the mutatee function in which the snippet is inserted. |
| <code>default</code> | The default domain (domain not specified) is the domain of snippet variables. |

Table 3: DynC API Domains

Example:

```
inf'printf("n is equal to %d.\n", ++global'n);
```

This would increment and print the value of the mutatee global variable n.

2.2 Control Flow

2.2.1 Comments

Block and line comments work as they do in C or C++.

Example:

```
/*  
 * This is a comment.  
 */  
4 int i; // So is this.
```

2.2.2 Conditionals

Use `if` to conditionally execute code. Example:

```
1 if(x == 0){  
    inf'printf("x == 0.\n");  
}
```

The `else` command can be used to specify code executed if a condition is not true. Example:

```
2 if(x == 0){  
    inf'printf("x == 0.\n");  
}else if(x > 3){  
    inf'printf("x > 3.\n");  
}else{  
7    inf'printf("x < 3 but x != 0.\n");  
}
```

2.2.3 First-Only Code Block

Code enclosed by a pair of `% <code> %` is executed only once by a snippet. First-only code blocks can be useful for declaring and initializing variables, or for any task that needs to be executed only once. Any number of first-only code blocks can be used in a dynC code snippet.

A first-only code block is equivalent to the following:

```
static int firstTime = 0;  
if(firstTime == 0){  
3    <code>  
    firstTime = 1;  
}
```

DynC will only execute the code in a first-only section the first time a snippet is executed. If `createSnippet(...)` is called multiple times and is passed the same name, then the first-only code will be executed only once: the first time that any of those snippets with the same name is executed. In contrast, if a snippet is created by calling `createSnippet(...)` with a unique snippet name (or if a name is unspecified), the first-only code will be executed only once upon reaching the first point encountered in the execution of the mutatee where the returned `BPatch_Snippet` is inserted.

Example Touch:

```
{%
  inf'printf("Function %s has been touched.\n", dyninst'function_name);
%}
```

If `createSnippet(...)` is passed the code in Example Touch and the name "fooTouchSnip" and the returned `BPatch_snippet` is inserted at the entry to function `foo`, the output would be:

```
2   Function foo has been touched.
   (mutatee exit)
```

If the dynC code in Example Touch is passed to `createSnippet(...)` multiple times and each snippet is given the same name, but is inserted at the entries of the functions `foo`, `bar`, and `run` respectively, the output would be:

```
   Function foo has been touched.
   (mutatee exit)
```

Creating the snippets with distinct names (e.g. `createSnippet(...)` is called with the dynC code in Example Touch multiple times and the snippets are named "fooTouchSnip", "barTouchSnip", "runTouchSnip") would produce an output like:

```
3   Function foo has been touched.
   Function bar has been touched.
   Function run has been touched.
   (mutatee exit)
```

A cautionary note: the use of first-only blocks can be expensive, as a conditional must be evaluated each time the snippet is executed. If the option is available, one may opt to insert a dynC snippet initializing all global variables at the entry point of `main`.

2.3 Variables

DynC allows for the creation of *snippet local* variables. These variables are in scope only within the snippet in which they are created.

For example,

```
1   int i;
   i = 5;
```

would create an uninitialized variable named `i` of type integer. The value of `i` is then set to 5. This is equivalent to:

```
int i = 5;
```

2.3.1 Static Variables

Every time a snippet is executed, non-static variables are reinitialized. To create a variable with value that persists across executions of snippets, declare the variable as static.

Example:

```
int i = 10;
inf'printf("i is %d.\n", i++);
```

If the above is inserted at the entrance to a function that is called four times, the output would be:

```
i is 10.
i is 10.
3 i is 10.
i is 10.
```

Adding `static` to the variable declaration would make the value of `i` persist across executions:

```
1 static int i = 10;
inf'printf("i is %d.\n", i++);
```

Produces:

```
i is 10.
i is 11.
3 i is 12.
i is 13.
```

A variable declared in a first-only section will also behave statically, as the initialization occurs only once.

```
1 {%
int i = 10;
%}
```

2.3.2 An Explanation of the Internal Workings of DynC Variable Creation

DynC uses the DyninstAPI function `malloc(...)` to allocate dynC declared variables when `createSnippet(...)` is called. The variable name is mangled with the name of the snippet passed to `createSnippet`. Thus, variables declared in dynC snippets are accessible only to those snippets created by calling `createSnippet(...)` with the same name.

If the variables are explicitly initialized, dynC sets the value of the variable with a `BPatch_arithExpr(BPatch_assign...)` snippet. Because of this, each time the snippet is executed, the value is reset to the initialized value. If, however the variables are not explicitly initialized, they are automatically set to a type-specific zero-value. Scalar variables are set to 0, and c-strings are set to empty, null-terminated strings (i.e. "").

If a variable is declared with the `static` keyword, then the initialization is performed as if in a first-only block (see section 2.2.3). Thus, a variable is initialized only the first time that snippet is executed, and subsequent executions of the variable initialization are ignored.

2.3.3 Creating Global Variables That Work With DynC

To declare a global variable that is accessible to all snippets inserted into a mutatee, one must use the DyninstAPI `BPatch_addressSpace::malloc(...)` method (see [Dyninst Programmer's Guide](#)). This code is located in mutator code (*not* in dynC code).

myMutator.C:

```
...
2 // Creates a global variable of type in named globalIntN
  myAddressSpace->malloc(myImage->getType("int"), "globalIntN");

  // file1 and file2 are FILE *, entryPoint and exitPoint are BPatch_point
  BPatch_snippet *snippet1 = dynC::createSnippet(file1, &entryPoint, "mySnippet1");
7  BPatch_snippet *snippet2 = dynC::createSnippet(file2, &exitPoint, "mySnippet2");

  assert(snippet1);
  assert(snippet2);

12  myAdressSpace->insertSnippet(snippet1, &entryPoint);
  myAdressSpace->insertSnippet(snippet2, &exitPoint);

  // run the mutatee
  ((BPatch_process *)myAdressSpace)->continueExecution();
17  ...

file1:

  {%
    global'globalIntN = 0; // initialize global variable in first-only section
3  %}
  inf'printf("Welcome to function %s. Global variable globalIntN = %d.\n",
    dyninst'function_name, global'globalIntN++);

file2:

  inf'printf("Goodbye from function %s. Global variable globalIntN = %d.\n",
    dyninst'function_name, global'globalIntN++);
```

When run, the output from the instrumentation would be:

```
3  Welcome to function foo. Global variable globalIntN = 0.
  Goodbye from function foo. Global variable globalIntN = 1.
  Welcome to function foo. Global variable globalIntN = 2.
  Goodbye from function foo. Global variable globalIntN = 3.
  Welcome to function foo. Global variable globalIntN = 4.
  Goodbye from function foo. Global variable globalIntN = 5.
```

2.3.4 Data Types

DynC supported data types are restricted by those supported by Dyninst: `int`, `long`, `char *`, and `void *`. Integer and c-string primitives are also recognized:

Example:

```
int i = 12;
char *s = "hello";
```

2.3.5 Pointers

Pointers are dereferenced with the prefix `*<variable>` and the address of variable is specified by `&<variable>`. For example, in reference to the previous example from section 2.3.4, the statement `*s` would evaluate to the character `h`.

2.3.6 Arrays

Arrays in DynC behave much the same way they do in C.

Example:

```
int array[3] = {1, 2, 3};
char *names[] = {"Mark", "Phil", "Deb", "Tracy"};
names[2] = "Gwen" // change Deb to Gwen
inf'printf("The seventh element of mutArray is %d.\n", global'mutArray[6]); //Mutatee array
if(inf'strcmp(*names, "Mark") == 0){} // This will evaluate to true.
```

2.4 DynC Limitations

The DynC, while quite expressive, is limited to those actions supported by the DyninstAPI. As such, it lacks certain abilities that many programmers have come to expect. These differences will be discussed in an exploration of those C abilities that dynC lacks.

2.4.1 Loops

There are no looping structures in DynC.

2.4.2 Enums, Unions, Structures

These features present a unique implementation challenge and are in development. Look to future revisions for full support for enums, unions, and structures.

2.4.3 Preprocessing

DynC does not allow C-style preprocessing macros or importation. Rather than `#define` statements, constant variables are recommended.

2.4.4 Functions

Specifying functions is beyond the scope of the DynC language. DyninstAPI has methods for dynamically loading code into a mutatee, and these loaded functions can be used in DynC snippets.

A The Dyninst Domain

The dyninst domain has quite a few useful values and functions:

| Identifier | Type | Where Valid | Description |
|--------------------------------|---------------------|--------------------------|---|
| <code>function_name</code> | <code>char *</code> | Within a function | Evaluates to the name of the current function. Call to <code>createSnippet(...)</code> must specify a <code>BPatch_point</code> . |
| <code>module_name</code> | <code>char *</code> | Anywhere | Evaluates to the name of the current module. Call to <code>createSnippet(...)</code> must specify a <code>BPatch_point</code> . |
| <code>bytes_accessed</code> | <code>int</code> | At a memory operation | Evaluates to the number of bytes accessed by a memory operation. |
| <code>effective_address</code> | <code>void *</code> | At a memory operation | Evaluates the effective address of a memory operation. |
| <code>original_address</code> | <code>void *</code> | Anywhere | Evaluates to the original address where the snippet was inserted. |
| <code>actual_address</code> | <code>void *</code> | Anywhere | Evaluates to the actual address of the instrumentation. |
| <code>return_value</code> | <code>void *</code> | Function exit | Evaluates to the return value of a function. |
| <code>thread_index</code> | <code>int</code> | Anywhere | Returns the index of the thread the snippet is executing on. |
| <code>tid</code> | <code>int</code> | Anywhere | Returns the id of the thread the snippet is executing on. |
| <code>dynamic_target</code> | <code>void *</code> | At calls, jumps, returns | Calculates the target of a control flow instruction. |
| <code>break()</code> | <code>void</code> | Anywhere | Causes the mutatee to execute a breakpoint. |
| <code>stopthread()</code> | <code>void</code> | Anywhere | Stops the thread on which the snippet is executing. |

Table 4: Dyninst Domain Values