




Abstract

On the quest of High Performance Computing systems capable of achieving high parallelism on scientific applications, a trend of adding acceleration devices with specialized hardware architectures can be observed. An evidence of this are the 102 supercomputers of the latest top 500 list that use an accelerator or co-processor of which 86 use GPGPUs. Furthermore, new programming languages and frameworks have been created to support and program accelerators. While each vendor may offer its own programming solution, traditional programming frameworks such as OpenMP works towards solving this issue by adapting its standard starting at version 4.0 to support target devices. Although currently it is not the only option, all solutions share a common objective of increasing code portability across vendors as well as multiple architectural generations.

However, maturity and correctness of these solution needs to be tested. In particular for OpenMP, there is a breach between the moment the standard is released and when compilers and vendors implement support for such standard. Hence, there is a need of a methodology to assess the quality of an implementation of the OpenMP standard, as well as evaluate system compliance during system deployments. Such methodology should reveal possible bugs or conflicts on the interpretation of the standard, as well as provide metrics to characterize the quality level of many implementation.

To this end, this work presents a tests suite for OpenMP 4.5 that is currently under development. This suite follows a well thought methodology that was established in advanced to provide a common and steady ground, allowing vendors, systems developers and programmers to asses the level of readiness and maturity of an OpenMP 4.5 implementation for a particular vendor.


Methodology



There are three possible positive outcomes of our methodology to build tests. Either a test passes through all the checks and makes it to the validation suite, or it uncovers a bug in the vendor implementation of the OpenMP 4.5 standard, or it highlights a contentious concept or text that is easy to misinterpret which needs to be brought to the attention of the OpenMP community and specification developers. All tests are written agnostic to where they are executed (host vs. target). After a test executes the output indicates if the test passed or failed and where it was executed (host or target).

OpenMP offloading

Host centric execution of code: Offloading directives provides the compiler with hints to create device executable code, as well as inline all the necessary calls for device initialization, code execution and data movement between host and device. OpenMP frees the programmer from bookkeeping data allocation and movement, as well as separate compilation of code for host and device. OpenMP 4.5 in particular provides more control to the programmer to handle data movement between host and device.



```
#pragma omp target map(tofrom: myVar) if(myCondition) device(2) depend(out: a)
{
  myVar++;
}
```

Simple Test Cases

```
int test_map_device() {
  int sum_dev = omp_get_num_devices(), num(num_dev), errors = 0;
  int h_matrix = (int*) malloc(num_dev * sizeof(int));
  for (int dev = 0; dev < num_dev; ++dev)
    h_matrix[dev*N] = dev;
  #pragma omp target map(from: h_matrix[dev*N:N]) device(dev)
  #pragma omp target map(from: h_matrix[dev*N:N]) device(dev)
  for (int i = 0; i < N; ++i)
    h_matrix[dev*N + i] = dev;
  // end target data
}
// checking results
errors = 0;
for (int dev = 0; dev < num_dev; ++dev) {
  sum_dev += h_matrix[dev*N];
  for (int i = 1; i < N; ++i)
    sum_dev += h_matrix[dev*N + i];
  errors += (dev * N != sum_dev);
}
return errors;
}
```

Offloading Multiple devices: Distribute each row of a matrix to one of the available devices (lines 5-14). Each iteration performs data movement and computation in a different device. Target data region maps a portion of the matrix to each device (line 6). Computation is performed on the target region (line 8).

```
class B {
public:
  static double VAR;
  B() {}
  static void modify(double, res) { ... }
  #pragma omp target map(tofrom: var)
  { *res = B::VAR; }
};
double B::VAR = 1.0;
int test_static() {
  int errors = 0;
  double exp = 1.0, res = 0.0;
  B::modify(8.0, res);
  errors += res != exp;
  return errors;
}
```

Mapping static attribute of a class: The unique value of the static VAR is mapped inside the method of a class with a target region and the map clause (lines 6-9). An OpenMP 4.5 capable compiler should capture the static variable (VAR) and map it to and from the device.

```
double sum = 0.0;
double h_array = (double*) malloc(N * sizeof(double));
double in_1 = (double*) malloc(N * sizeof(double));
double in_2 = (double*) malloc(N * sizeof(double));
#pragma omp taskwait
{ for (int i = 0; i < N; ++i)
  in_1[i] = 1; }
#pragma omp taskwait
{ for (int i = 0; i < N; ++i)
  in_2[i] = 2; }
#pragma omp target enter data nowait
{ map(tofrom: h_array[0:N]) map(to: in_1[0:N]) \
  map(to: in_2[0:N]) depend(inout: h_array) \
  depend(in: in_1) depend(in: in_2) }
#pragma omp target exit data nowait
{ h_array[i] = in_1[i] * in_2[i]; }
#pragma omp target enter data nowait
{ map(tofrom: h_array[0:N]) depend(inout: h_array) }
#pragma omp taskwait
{ sum = h_array[0]; }
#pragma omp taskwait
errors = 2.0 * N != sum;
```

Task dependencies: Task graph made out of host and target tasks with in, out and inout dependencies. Asynchronous behavior specified using the *nowait* clause. Data mapping tasks are separated from computation tasks. Signals between host and target in both directions

Complex Test Cases

```
typedef struct node {
  double data;
  struct node *next;
} node;
void map_ll(node * head) {
  if (!head) return;
  #pragma omp target enter data map(to:head[0])
  while(head->next) {
    Note: explicit attachment
    node * cur = head->next;
    #pragma omp target enter data map(to:cur[0])
    #pragma omp target
    head->next = cur;
  }
}
void unmap_ll(node * head) {
  if (!head) return;
  #pragma omp target exit data map(from:head[0].data)
  while(head->next) {
    Note: only copies back the data element to avoid
    overwriting next pointer
    #pragma omp target exit data map(from:head[0].next[0].data)
  }
}
```

Mapping Linked list to device: The *map_ll* function (line 5) uses *target enter data* directive to first map the head of the linked list, and then map the pointer to the next link using array dereferencing syntax. The *unmap_ll* (line 19) function explicitly copies the data using *map-type* from with *target exit data map*.

```
#define RealType double
#define N 10
#pragma omp declare target
class MyVector {
public:
  inline RealType operator()(int i, int j, int k) const {
    return N*i*Length[2]-(j)-Length[1]+i; }
  inline RealType operator()(int i, int j, int k) {
    return N*i*Length[2]-(j)-Length[1]+i; }
  MyVector(int i, int m, int n) {
    Length[0] = i;
    Length[1] = m;
    Length[2] = n;
    X = new RealType[1-m*n];
  }
  RealType *getdata() const { return X; }
  RealType *getdata() const { return X; }
  int getSize() const { return Length[0]*Length[1]*Length[2]; }
  RealType *X;
};
#pragma omp end declare target
int main() {
  MyVector gamma(N, N, N);
  int size = gamma.getSize();
  #pragma omp target enter data map(to:gamma)
  #pragma omp target enter data map(to:gamma.X[0:size]) map(to:gamma.Length)
  #pragma omp target
  for (int i = 0; i < N; ++i)
    for (int j = 0; j < N; ++j)
      for (int k = 0; k < N; ++k)
        gamma(i, j, k) = i * k;
  #pragma omp target exit data map(from:gamma.X[0:size])
  for (int i = 0; i < N; ++i)
    for (int j = 0; j < N; ++j)
      for (int k = 0; k < N; ++k)
        cout << gamma(i, j, k) << " ";
  return 0;
}
```

Deep Copy of Classes: This code came from analyzing a full scale ECP application. It uses the *declare target* directive (line 4 to 36) to ensures that procedures and global variables can be executed and data can be accessed on the device. When the C++ methods are encountered, device-specific versions of the routines are created that can be called from a target region. Deep copy is performed through the use of *target enter data* (lines 43 and 44) by first mapping the class and then the individual class members. Computation is performed on the device (line 46). After computation is over, the data is copy back to the host (line 52).

Current Snapshot

Our intention is to develop a test suite for the entire OpenMP 4.5 specification. We classify our tests by directives and clauses. The aforementioned methodology was used. Some of them have resulted in multiple bug reports. Following is a snapshot of the current suite

- Tested using 4 different compilers: Clang, IBM XL, GCC and Cray CCE.
- Target platforms:
 - Titan Cray XK7:** AMD Opteron x64 + Nvidia K20X
 - SummitDev:** IBM Power8 + NVIDIA TESLA P100

Test	Lang	Compiler: Clang 3.8.0 ORNL Version	Platform: Summitdev	Platform: Titan	Platform: Summitdev
offloading_success.c	C	PASS	PASS	PASS	PASS
offloading_success.cpp	C++	PASS	PASS	PASS	PASS
targettest_target_is_device_ptr.c	C	PASS	PASS	PASS	PASS
targettest_target_map_array_default.c	C	PASS	PASS	PASS	PASS
targettest_target_map_pointer.c	C	PASS	PASS	PASS	PASS
targettest_target_map_pointer_default.c	C	PASS	PASS	PASS	PASS
targettest_target_map_scalar_default.c	C	PASS	PASS	PASS	PASS
targettest_target_if.c	C	PASS	PASS	PASS	PASS
targettest_target_map_global_arrays.c	C	PASS	PASS	PASS	PASS
targettest_target_map_local_array.c	C	PASS	PASS	PASS	PASS
target_data/test_target_data_if.c	C	PASS	PASS	PASS	PASS
target_data/test_target_data_map.c	C	PASS	PASS	PASS	PASS
target_enter_exit_data/test_target_enter_exit_data_async.c	C	PASS	PASS	PASS	PASS
target_enter_exit_data/test_target_enter_exit_data_map_global_array.c	C	PASS	PASS	PASS	PASS
target_enter_exit_data/test_target_enter_exit_data_map_malloced_array.c	C	PASS	PASS	C.F.AIL	PASS
targettest_target_map_struct_default.c	C	PASS	PASS	C.F.AIL	PASS
targettest_target_device.c	C	PASS	PASS	C.F.AIL	PASS
targettest_target_async.c	C	PASS	R.F.AIL	PASS	PASS
target_data/test_target_data_map_array_sections.c	C	PASS	PASS	R.F.AIL	PASS
target_data/test_target_data_map_classes.cpp	C++	PASS	PASS	C.F.AIL	PASS
target_data/test_target_data_map_devices.c	C	PASS	PASS	C.F.AIL	PASS
target_enter_exit_data/test_target_enter_exit_data_devices.c	C	PASS	PASS	C.F.AIL	PASS
target_enter_exit_data/test_target_enter_exit_data_struct.c	C	PASS	PASS	C.F.AIL	PASS
application_kernels/linked_list.c	C	PASS	PASS	C.F.AIL	PASS
target_data/test_target_data_use_device_ptr.c	C	PASS	PASS	R.F.AIL	C.F.AIL
targettest_target_defaultmap.c	C	PASS	R.F.AIL	R.F.AIL	PASS
targettest_target_map_classes_default.cpp	C++	R.F.AIL	R.F.AIL	C.F.AIL	C.F.AIL
target_enter_exit_data/test_target_enter_exit_data_classes.cpp	C++	C.F.AIL	R.F.AIL	C.F.AIL	C.F.AIL