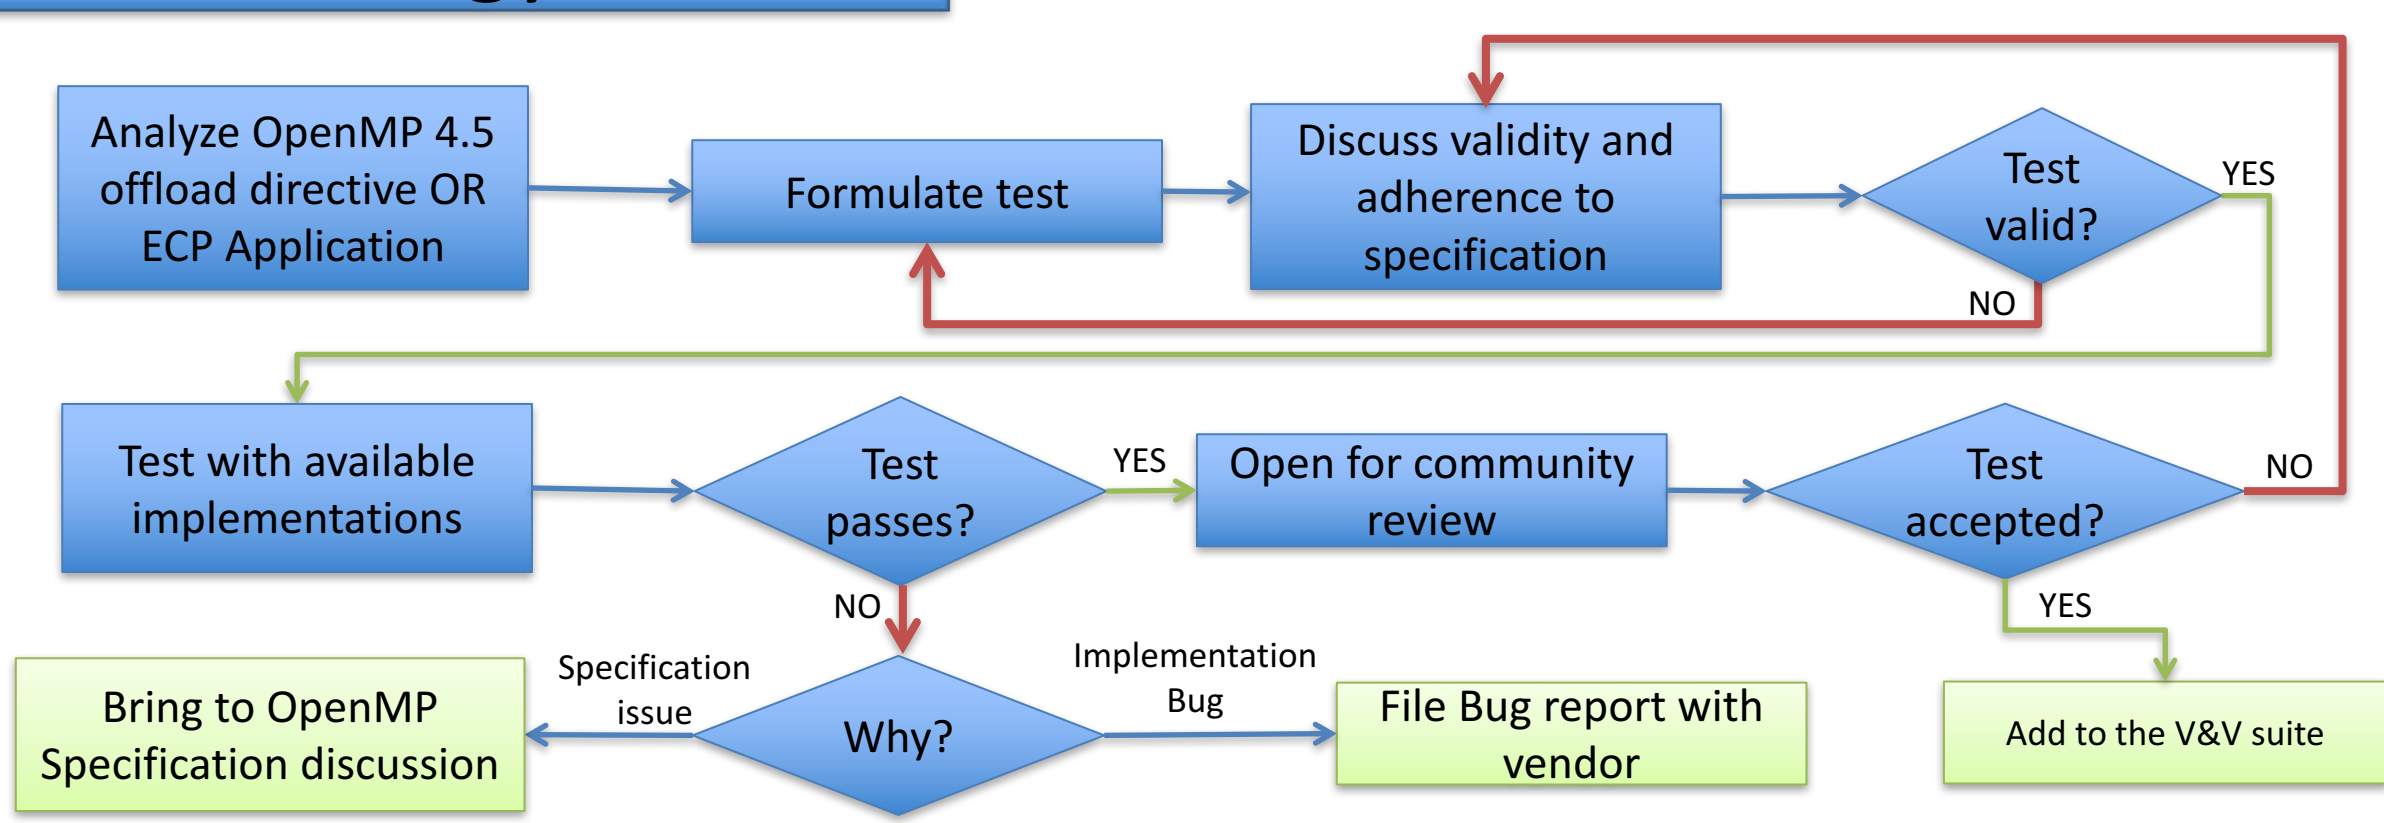


Abstract

- Building Validation and Verification Suite focusing on the offloading features of OpenMP 4.5 and beyond
- Creating functional and unit tests
- Collaboration with ECP application developers to create representative scientific use cases
- Project in collaboration with OpenMP community, Oak Ridge National Lab (ORNL), Argonne National Lab (ANL), University of Delaware (UDEL), IBM, LLVM, Cray, GCC
- Project uncovers compiler and runtime implementation bugs and ambiguities in the OpenMP 4.5 specification
- Compilers that this project use for evaluation include Clang/LLVM, GNU GCC, IBM XL, Cray CCE
- Target platforms include ORNL's TITAN and SummitDev (representative exascale system)
- Project is open for collaboration
- Feedback and suggestions from the HPC community is welcome - Contact any of the authors

Methodology



There are three possible positive outcomes of the process we have adopted. Either a test passes through all the checks and makes it to the validation suite, uncovers a bug in the vendor implementation of the OpenMP 4.5 standard, or highlights a contentious concept or text that is easy to misinterpret and brings it 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)
{
  myVar++;
}
```

Simple Test Cases

```
int test_map_device() {
  int num_dev = omp_get_num_devices(), num_mmm_dev, errors = 0;
  int h_matrix = (int*) malloc(num_dev * num_mmm_dev * sizeof(int));
  for (int dev = 0; dev < num_dev; ++dev) {
    #pragma omp target data 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
    } // end target data
  }
  // checking results
  errors = 0;
  for (int dev = 0; dev < num_dev; ++dev) {
    int dev_err = 0;
    for (int i = 0; i < N; ++i)
      errors += h_matrix[dev*N + i];
    errors -= (dev * N * num_mmm_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: res(0:1))
    {
      res = B::VAR;
    }
  }
  double B::VAR = 1.0;
};

int test_static() {
  double exp = 1.0, res = 0.0;
  B::modify(&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 mm = 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 task depend(out: in_1) shared(in_1)
{ for (int i = 0; i < N; ++i)
  in_1[i] = 1; }

#pragma omp task depend(out: in_2) shared(in_2)
{ for (int i = 0; i < N; ++i)
  in_2[i] = 2; }

#pragma omp target enter data nowait \
  map(to: h_array[0:N]) map(to: in_1[0:N]) \
  map(to: in_2[0:N]) depend(out: h_array) \
  depend(in: in_1) depend(in: in_2)
{
  #pragma omp target enter data nowait \
  {
    h_array[i] = in_1[i] + in_2[i];
  }
}

#pragma omp target exit data nowait \
  map(from: h_array[0:N]) depend(inout: h_array)
{
  shared(mm, h_array);
  for (int i = 0; i < N; ++i)
    mm += h_array[i];
}

#pragma omp taskwait
errors = 2.0 * N != mm;
```

Task dependencies: Task graph composed of host and target tasks that have in and out dependencies between each other. Asynchronous behavior is specified using the *nowait* clause. Data map tasks are separated from computation tasks.

Complex Test Cases

```
typedef struct node {
  double data;
  struct node * next;
} node_t;

void map_ll(node_t * head) {
  if (!head) return;
  #pragma omp target enter data map(to:head[0] data)
  while(head->next) {
    Note: explicit attachment
    node_t * cur = head->next;
    #pragma omp target enter data map(to:cur[0] data)
    #pragma omp target
    {
      head->next = cur;
    }
  }
}

void unmap_ll(node_t * head) {
  if (!head) return;
  #pragma omp target enter 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 X[k-Length[2]-j+Length[1]-i]; }
  inline RealType operator()(int i, int j, int k)
  { return X[k-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[i*m*n];
  }
  RealType & getData() { return X; }
  RealType * getData() const { return X; }
  int getSize() const { return Length[0]*Length[1]*Length[2]; }
  int Length[3];
  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) = 1.0;
  #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)
        return 0;
  cout << gamma(i,j,k) << "\n";
}
```

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 ensure 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 divide our tests by directives. These tests have gone through the methodology described previously. Some of them have result in different 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	Platform: Summitdev		Platform: Titan	
		Compiler: Clang 3.8.0 ORNL Version	Compiler: IBM XL 13.1.6	Compiler: CCE 8.6.3	Compiler: GCC 7.1.1
offloading_success.c	C	PASS	PASS	PASS	PASS
offloading_success.cpp	C++	PASS	PASS	PASS	PASS
target/test_target_is_device_ptr.c	C	PASS	PASS	PASS	PASS
target/test_target_map_array_default.c	C	PASS	PASS	PASS	PASS
target/test_target_map_pointer.c	C	PASS	PASS	PASS	PASS
target/test_target_map_pointer_default.c	C	PASS	PASS	PASS	PASS
target/test_target_map_scalar_default.c	C	PASS	PASS	PASS	PASS
target/test_target_if.c	C	PASS	PASS	PASS	PASS
target/test_target_map_global_arrays.c	C	PASS	PASS	PASS	PASS
target/test_target_map_local_arrays.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 malloced array.c	C	PASS	PASS	PASS	PASS
target enter: exit data/test_target enter exit data map default.c	C	PASS	PASS	C.FAIL	PASS
target/test_target_device.c	C	PASS	PASS	C.FAIL	PASS
target/test_target_async.c	C	PASS	R.FAIL	PASS	PASS
target_data/test_target_data_map_array_sections.c	C	PASS	PASS	R.FAIL	PASS
target_data/test_target_data_map_classes.cpp	C++	PASS	PASS	C.FAIL	PASS
target_data/test_target_data_map_devices.c	C	PASS	PASS	C.FAIL	PASS
target enter: exit data/test_target enter exit data devices.c	C	PASS	PASS	C.FAIL	PASS
target enter: exit data/test_target enter exit data struct.c	C	PASS	PASS	C.FAIL	PASS
application_kernels/linked_list.c	C	PASS	PASS	C.FAIL	PASS
target_data/test_target_data_use_device_ptr.c	C	PASS	PASS	R.FAIL	C.FAIL
target/test_target_defaultmap.c	C	PASS	R.FAIL	R.FAIL	PASS
target/test_target_map_classes_default.cpp	C++	R.FAIL	R.FAIL	C.FAIL	C.FAIL
target enter: exit data/test_target enter exit_data_classes.cpp	C++	C.FAIL	R.FAIL	C.FAIL	C.FAIL