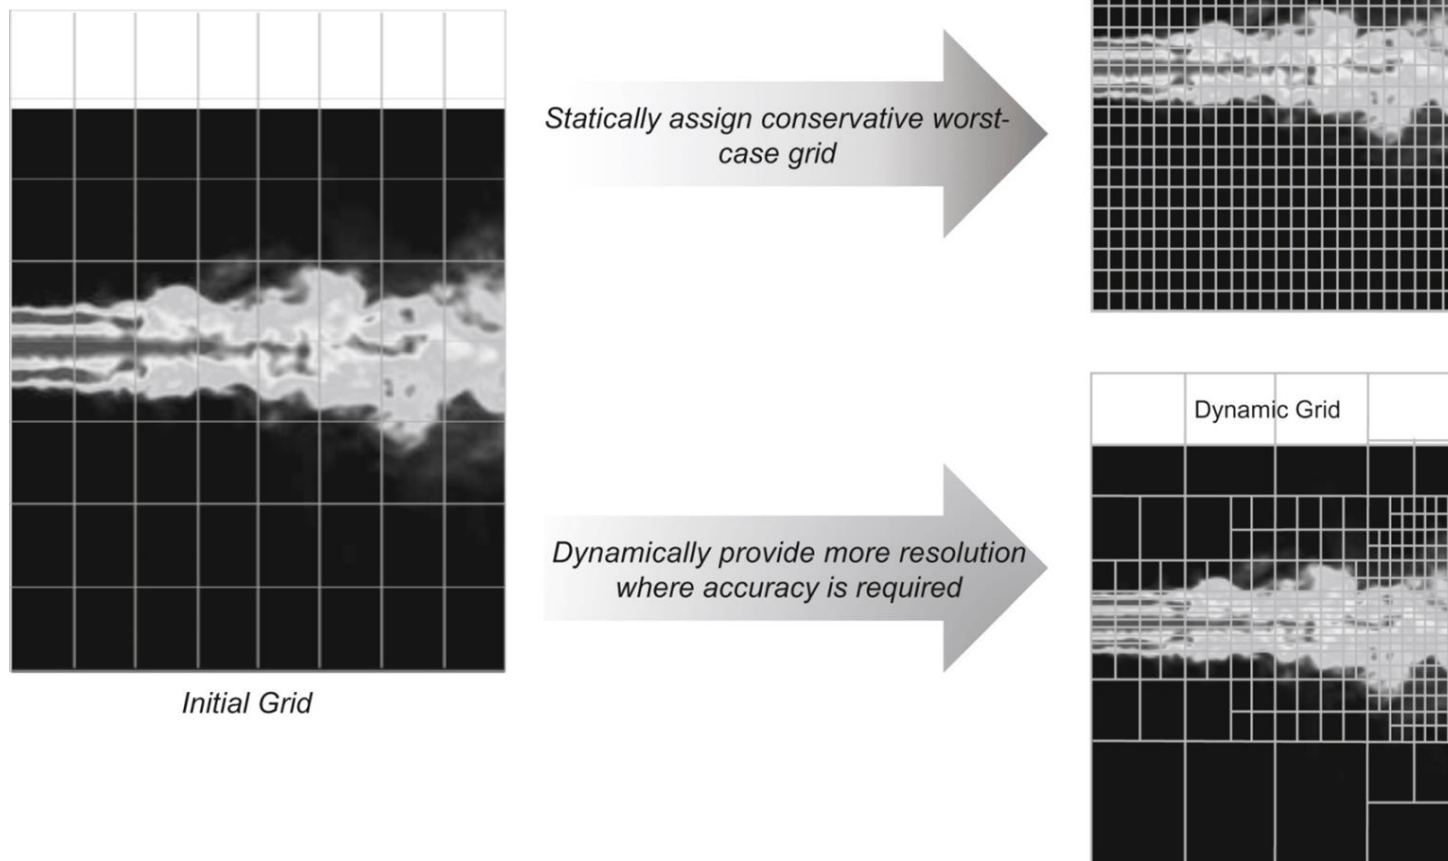




# Chapter-13

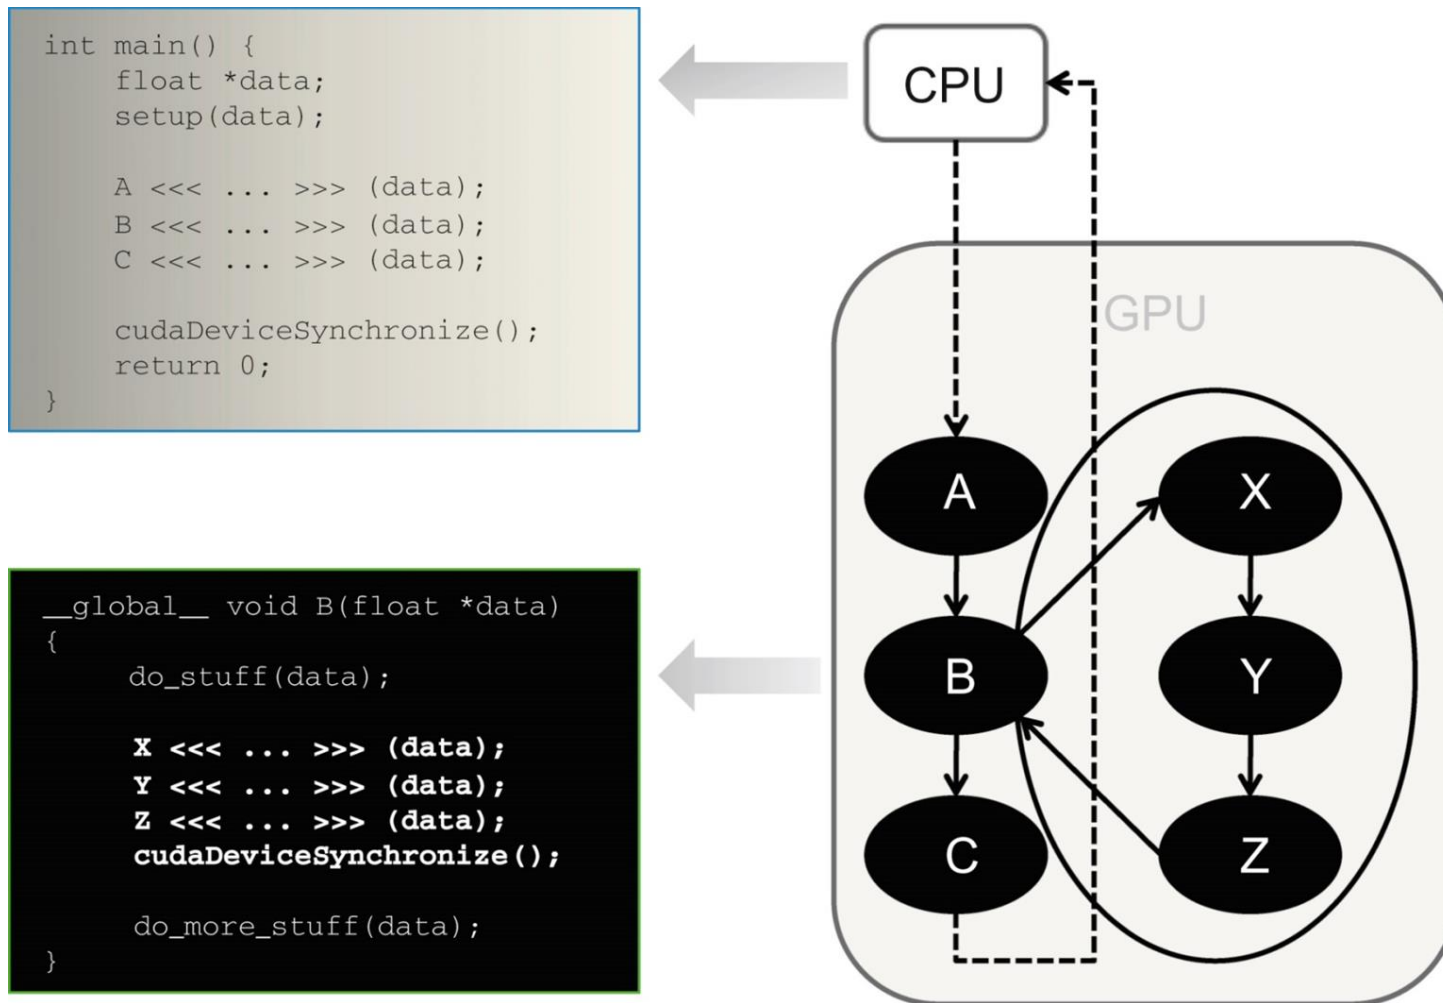
## CUDA dynamic parallelism



**FIGURE 13.1:** Fixed versus dynamic grids for a turbulence simulation model.



**FIGURE 13.2:** Kernel launch patterns for algorithms with dynamic work variation, with and without dynamic parallelism.



**FIGURE 13.3:** A simple example of a kernel (B) launching three kernels (X, Y, and Z).

```
01  __global__ void kernel(unsigned int* start, unsigned int* end, float* someData,  
02      float* moreData) {  
03  
04      unsigned int i = blockIdx.x*blockDim.x + threadIdx.x;  
05      doSomeWork(someData[i]);  
06  
07      for(unsigned int j = start[i]; j < end[i]; ++j) {  
08          doMoreWork(moreData[j]);  
09      }  
10  
11  }
```

**FIGURE 13.4:** A simple example of a hypothetical parallel algorithm coded in CUDA without dynamic parallelism.

```

01  __global__ void kernel_parent(unsigned int* start, unsigned int* end,
02      float* someData, float* moreData) {
03
04      unsigned int i = blockIdx.x*blockDim.x + threadIdx.x;
05      doSomeWork(someData[i]);
06
07      kernel_child <<< ceil((end[i]-start[i])/256.0) , 256 >>>
08          (start[i], end[i], moreData);
09
10  }
11
12  __global__ void kernel_child(unsigned int start, unsigned int end,
13      float* moreData) {
14
15      unsigned int j = start + blockIdx.x*blockDim.x + threadIdx.x;
16
17      if(j < end) {
18          doMoreWork(moreData[j]);
19      }
20
21  }

```

**FIGURE 13.5:** A revised example using CUDA dynamic parallelism.

```
__device__ int value;
__device__ void x() {
    value = 5;
    child<<< 1, 1 >>>(&value);
}
```

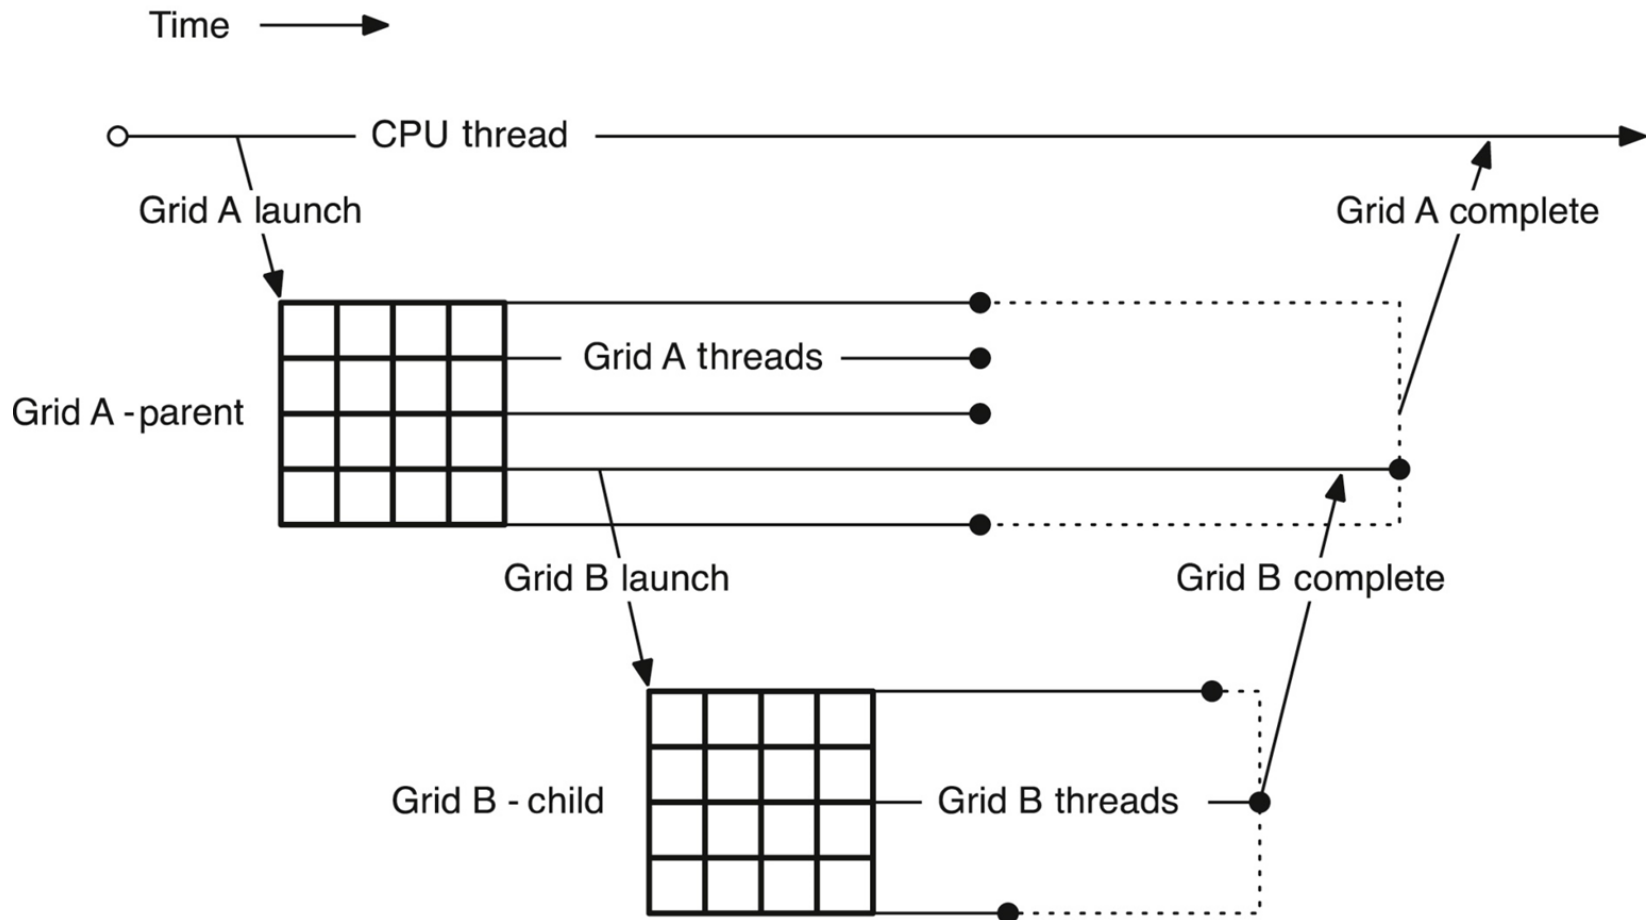
(A) Valid—"value" is global storage

```
__device__ void y() {
    int value = 5;
    child<<< 1, 1 >>>(&value);
}
```

(B) Invalid—"value" is local storage

**FIGURE 13.6:** Passing a pointer as an argument to a child kernel.

## Parent-child launch nesting



**FIGURE 13.7:** Completion sequence for parent and child grids.



```

01 #include <stdio.h>
02 #include <cuda.h>
03
04 #define MAX_TESS_POINTS 32
05
06 //A structure containing all parameters needed to tessellate a Bezier line
07 struct BezierLine {
08     float2 CP[3]; //Control points for the line
09     float2 vertexPos[MAX_TESS_POINTS]; //Vertex position array to tessellate into
10     int nVertices; //Number of tessellated vertices
11 };
12
13 __global__ void computeBezierLines(BezierLine *bLines, int nLines) {
14     int bidx = blockIdx.x;
15     if(bidx < nLines){
16         //Compute the curvature of the line
17         float curvature = computeCurvature(bLines);
18
19         //From the curvature, compute the number of tessellation points
20         int nTessPoints = min(max((int)(curvature*16.0f),4),32);
21         bLines[bidx].nVertices = nTessPoints;
22
23         //Loop through vertices to be tessellated, incrementing by blockDim.x
24         for(int inc = 0; inc < nTessPoints; inc += blockDim.x){
25             int idx = inc + threadIdx.x; //Compute a unique index for this point
26             if(idx < nTessPoints){
27                 float u = (float)idx/(float)(nTessPoints-1); //Compute u from idx
28                 float omu = 1.0f - u; //pre-compute one minus u
29                 float B3u[3]; //Compute quadratic Bezier coefficients
30                 B3u[0] = omu*omu;
31                 B3u[1] = 2.0f*u*omu;
32                 B3u[2] = u*u;
33                 float2 position = {0,0}; //Set position to zero
34                 for(int i = 0; i < 3; i++){
35                     //Add the contribution of the i'th control point to position
36                     position = position + B3u[i] * bLines[bidx].CP[i];
37                 }
38                 //Assign value of vertex position to the correct array element
39                 bLines[bidx].vertexPos[idx] = position;
40             }
41         }
42     }
43 }
44
45 #define N_LINES 256
46 #define BLOCK_DIM 32
47
48 int main( int argc, char **argv ) {
49     //Allocate and initialize array of lines in host memory
50     BezierLine *bLines_h = new BezierLine[N_LINES];
51     initializeBLines(bLines_h);
52
53     //Allocate device memory for array of Bezier lines
54     BezierLine *bLines_d;
55     cudaMalloc((void**)&bLines_d, N_LINES*sizeof(BezierLine));
56     cudaMemcpy(bLines_d, bLines_h, N_LINES*sizeof(BezierLine), cudaMemcpyHostToDevice);
57
58     //Call the kernel to tessellate the lines
59     computeBezierLines<<<N_LINES, BLOCK_DIM>>>>(bLines_d, N_LINES);
60
61     cudaFree(bLines_d); //Free the array of lines in device memory
62     delete[] bLines_h; //Free the array of lines in host memory
63 }

```

**FIGURE 13.8:** Bezier curve calculation without dynamic parallelism (support code in Fig. A13.8).

```

01 struct BezierLine {
02     float2 CP[3]; //Control points for the line
03     float2 *vertexPos; //Vertex position array to tessellate into
04     int nVertices; //Number of tessellated vertices
05 };
06 __global__ void computeBezierLines_parent(BezierLine *bLines, int nLines) {
07     //Compute a unique index for each Bezier line
08     int lidx = threadIdx.x + blockDim.x*blockIdx.x;
09     if(lidx < nLines){
10         //Compute the curvature of the line
11         float curvature = computeCurvature(bLines);
12
13         //From the curvature, compute the number of tessellation points
14         bLines[lidx].nVertices = min(max((int)(curvature*16.0f),4),MAX_TESS_POINTS);
15         cudaMalloc((void**)&bLines[lidx].vertexPos,
16                 bLines[lidx].nVertices*sizeof(float2));
17
18         //Call the child kernel to compute the tessellated points for each line
19         computeBezierLine_child<<<ceil((float)bLines[lidx].nVertices/32.0f), 32>>>
20             (lidx, bLines, bLines[lidx].nVertices);
21     }
22 }
23 __global__ void computeBezierLine_child(int lidx, BezierLine* bLines,
24     int nTessPoints) {
25     int idx = threadIdx.x + blockDim.x*blockIdx.x; //Compute idx unique to this vertex
26     if(idx < nTessPoints){
27         float u = (float)idx/(float)(nTessPoints-1); //Compute u from idx
28         float omu = 1.0f - u; //Pre-compute one minus u
29         float B3u[3]; //Compute quadratic Bezier coefficients
30         B3u[0] = omu*omu;
31         B3u[1] = 2.0f*u*omu;
32         B3u[2] = u*u;
33         float2 position = {0,0}; //Set position to zero
34         for(int i = 0; i < 3; i++) {
35             //Add the contribution of the i'th control point to position
36             position = position + B3u[i] * bLines[lidx].CP[i];
37         }
38         //Assign the value of the vertex position to the correct array element
39         bLines[lidx].vertexPos[idx] = position;
40     }
41 }
42 __global__ void freeVertexMem(BezierLine *bLines, int nLines) {
43     //Compute a unique index for each Bezier line
44     int lidx = threadIdx.x + blockDim.x*blockIdx.x;
45     if(lidx < nLines)
46         cudaFree(bLines[lidx].vertexPos); //Free the vertex memory for this line
47 }
48 int main( int argc, char **argv ) {
49     //Allocate array of lines in host memory
50     BezierLine *bLines_h = new BezierLine[N_LINES];
51     initializeBLines(bLines_h);
52
53     //Allocate device memory for array of Bezier lines
54     BezierLine *bLines_d;
55     cudaMalloc((void**)&bLines_d, N_LINES*sizeof(BezierLine));
56     cudaMemcpy(bLines_d,bLines_h, N_LINES*sizeof(BezierLine),cudaMemcpyHostToDevice);
57
58     computeBezierLines_parent<<<ceil((float)N_LINES/(float)BLOCK_DIM), BLOCK_DIM>>>
59         (bLines_d, N_LINES);
60
61     freeVertexMem <<<ceil((float)N_LINES/(float)BLOCK_DIM), BLOCK_DIM>>>
62         (bLines_d, N_LINES);
63     cudaFree(bLines_d); //Free the array of lines in device memory
64     delete[] bLines_h; //Free the array of lines in host memory
65 }

```

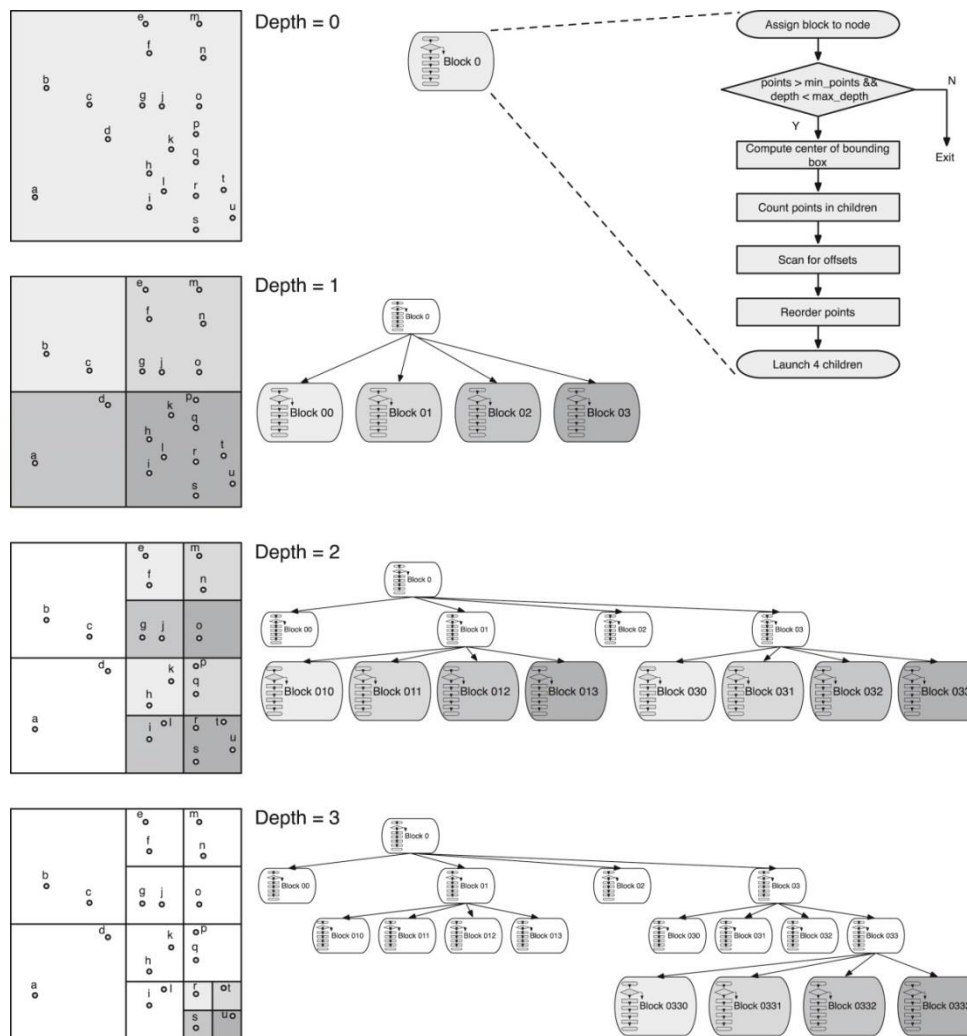
**FIGURE 13.9:** Bezier calculation with dynamic parallelism (support code in Fig. A13.8).

```
cudaStream_t stream;
// Create non-blocking stream
cudaStreamCreateWithFlags(&stream, cudaStreamNonBlocking);

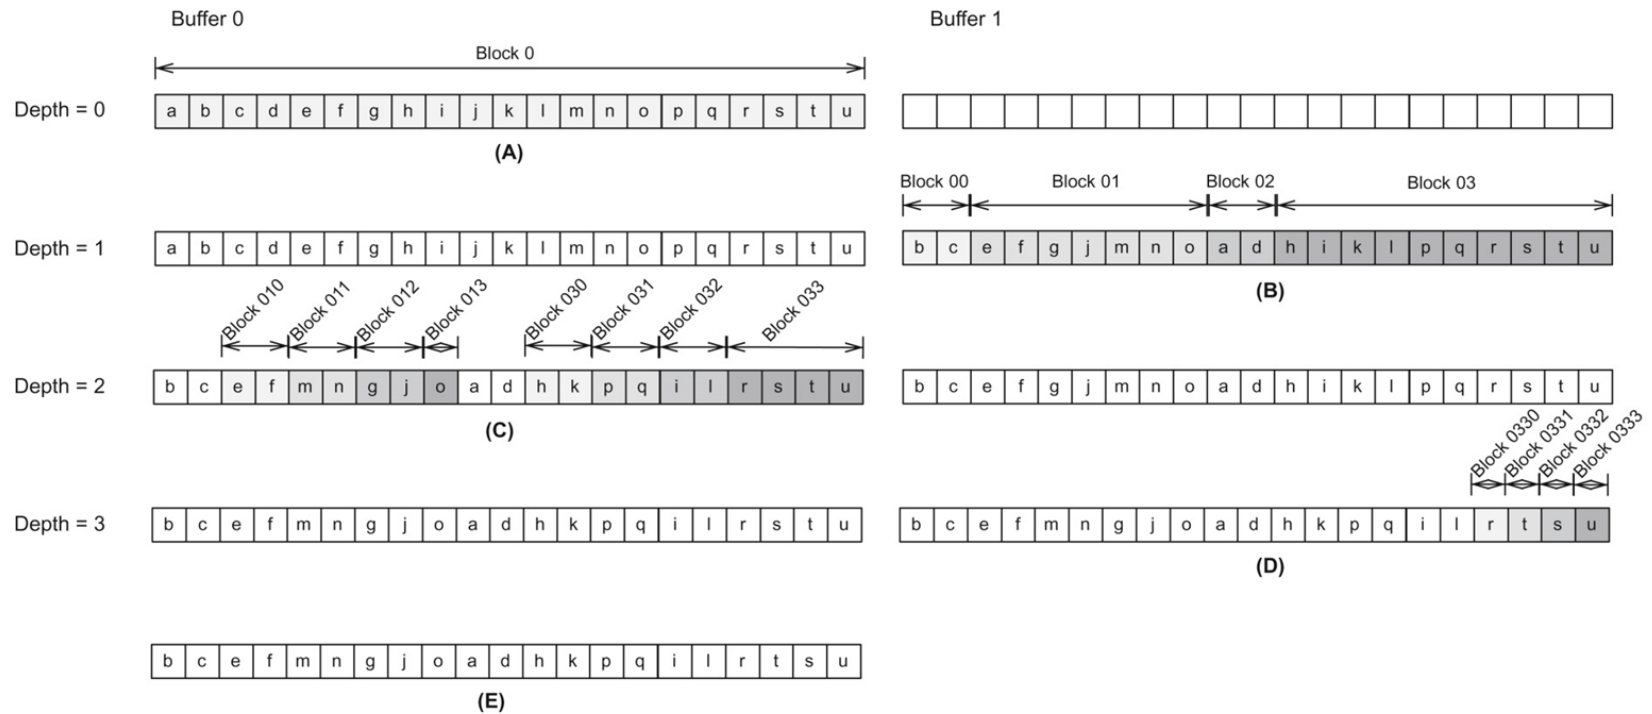
//Call the child kernel to compute the tessellated points for each line
computeBezierLine_child<<<ceil((float)bLines[lidx].nVertices/32.0f), 32, 0, stream>>>
    (lidx, bLines, bLines[lidx].nVertices);

// Destroy stream
cudaStreamDestroy(stream);
```

**FIGURE 13.10:** Child kernel launch with named streams.



**FIGURE 13.11:** Quadtree example. Each thread-block is assigned to one quadrant. If the number of points in a quadrant is more than 2, the block launches 4 child blocks. Shaded blocks are active blocks in each level of depth.



**FIGURE 13.12:** Quadtree example. At each level of depth, a block groups all points in the same quadrant together.

```

01  __global__ void build_quadtree_kernel
02      (Quadtree_node *nodes, Points *points, Parameters params) {
03      __shared__ int smem[8]; // To store the number of points in each quadrant
04
05      // The current node
06      Quadtree_node &node = nodes[blockIdx.x];
07      node.set_id(node.id() + blockIdx.x);
08      int num_points = node.num_points(); // The number of points in the node
09
10      // Check the number of points and its depth
11      bool exit = check_num_points_and_depth(node, points, num_points, params);
12      if(exit) return;
13
14      // Compute the center of the bounding box of the points
15      const Bounding_box &bbox = node.bounding_box();
16      float2 center;
17      bbox.compute_center(center);
18
19      // Range of points
20      int range_begin = node.points_begin();
21      int range_end   = node.points_end();
22      const Points &in_points = points[params.point_selector]; // Input points
23      Points &out_points = points[(params.point_selector+1) % 2]; // Output points
24
25      // Count the number of points in each child
26      count_points_in_children(in_points, smem, range_begin, range_end, center);
27
28      // Scan the quadrants' results to know the reordering offset
29      scan_for_offsets(node.points_begin(), smem);
30
31      // Move points
32      reorder_points(out_points, in_points, smem, range_begin, range_end, center);
33
34      // Launch new blocks
35      if (threadIdx.x == blockDim.x-1) {
36          // The children
37          Quadtree_node *children = &nodes[params.num_nodes_at_this_level];
38
39          // Prepare children launch
40          prepare_children(children, node, bbox, smem);
41
42          // Launch 4 children.
43          build_quadtree_kernel<<<4, blockDim.x, 8 *sizeof(int)>>>
44              (children, points, Parameters(params, true));
45      }
46  }

```

**FIGURE 13.13:** Quadtree with dynamic parallelism: recursive kernel (support code in Fig. A13.13).

```

001 // Check the number of points and its depth
002 __device__ bool check_num_points_and_depth(Quadtree_node &node, Points *points,
003                                           int num_points, Parameters params){
004     if(params.depth >= params.max_depth || num_points <= params.min_points_per_node) {
005         // Stop the recursion here. Make sure points[0] contains all the points
006         if(params.point_selector == 1) {
007             int it = node.points_begin(), end = node.points_end();
008             for (it += threadIdx.x ; it < end ; it += blockDim.x)
009                 if(it < end)
010                     points[0].set_point(it, points[1].get_point(it));
011         }
012         return true;
013     }
014     return false;
015 }
016
017 // Count the number of points in each quadrant
018 __device__ void count_points_in_children(const Points &in_points, int* smem,
019                                         int range_begin, int range_end, float2 center) {
020     // Initialize shared memory
021     if(threadIdx.x < 4) smem[threadIdx.x] = 0;
022     __syncthreads();
023     // Compute the number of points
024     for(int iter=range_begin+threadIdx.x; iter<range_end; iter+=blockDim.x){
025         float2 p = in_points.get_point(iter); // Load the coordinates of the point
026         if(p.x < center.x && p.y >= center.y)
027             atomicAdd(&smem[0], 1); // Top-left point?
028         if(p.x >= center.x && p.y >= center.y)
029             atomicAdd(&smem[1], 1); // Top-right point?
030         if(p.x < center.x && p.y < center.y)
031             atomicAdd(&smem[2], 1); // Bottom-left point?
032         if(p.x >= center.x && p.y < center.y)
033             atomicAdd(&smem[3], 1); // Bottom-right point?
034     }
035     __syncthreads();
036 }
037
038 // Scan quadrants' results to obtain reordering offset
039 __device__ void scan_for_offsets(int node_points_begin, int* smem){
040     int* smem2 = &smem[4];
041     if(threadIdx.x == 0){
042         for(int i = 0; i < 4; i++)
043             smem2[i] = i==0 ? 0 : smem2[i-1] + smem[i-1]; // Sequential scan
044         for(int i = 0; i < 4; i++)
045             smem2[i] += node_points_begin; // Global offset
046     }
047     __syncthreads();
048 }
049
050 // Reorder points in order to group the points in each quadrant
051 __device__ void reorder_points(
052     Points& out_points, const Points &in_points, int* smem,
053     int range_begin, int range_end, float2 center){
054     int* smem2 = &smem[4];
055     // Reorder points
056     for(int iter=range_begin+threadIdx.x; iter<range_end; iter+=blockDim.x){
057         int dest;
058         float2 p = in_points.get_point(iter); // Load the coordinates of the point
059         if(p.x<center.x && p.y>=center.y)
060             dest=atomicAdd(&smem2[0],1); // Top-left point?
061         if(p.x>=center.x && p.y>=center.y)
062             dest=atomicAdd(&smem2[1],1); // Top-right point?
063         if(p.x<center.x && p.y<center.y)
064             dest=atomicAdd(&smem2[2],1); // Bottom-left point?
065         if(p.x>=center.x && p.y<center.y)
066             dest=atomicAdd(&smem2[3],1); // Bottom-right point?
067         // Move point
068         out_points.set_point(dest, p);
069     }
070     __syncthreads();
071 }

```

**FIGURE 13.14:** Quadtree with dynamic parallelism: device functions (support code in Fig. A13.14).

```

072
073 // Prepare children launch
074 __device__ void prepare_children(Quadtree_node *children, Quadtree_node &node,
075                                 const Bounding_box &bbox, int *smem){
076     int child_offset = 4*node.id(); // The offsets of the children at their level
077
078     // Set IDs
079     children[child_offset+0].set_id(4*node.id()+ 0);
080     children[child_offset+1].set_id(4*node.id()+ 4);
081     children[child_offset+2].set_id(4*node.id()+ 8);
082     children[child_offset+3].set_id(4*node.id()+12);
083
084     // Points of the bounding-box
085     const float2 &p_min = bbox.get_min();
086     const float2 &p_max = bbox.get_max();
087
088     // Set the bounding boxes of the children
089     children[child_offset+0].set_bounding_box(
090         p_min.x , center.y, center.x, p_max.y); // Top-left
091     children[child_offset+1].set_bounding_box(
092         center.x, center.y, p_max.x , p_max.y); // Top-right
093     children[child_offset+2].set_bounding_box(
094         p_min.x , p_min.y , center.x, center.y); // Bottom-left
095     children[child_offset+3].set_bounding_box(
096         center.x, p_min.y , p_max.x , center.y); // Bottom-right
097
098     // Set the ranges of the children.
099     children[child_offset+0].set_range(node.points_begin(), smem[4 + 0]);
100     children[child_offset+1].set_range(smem[4 + 0], smem[4 + 1]);
101     children[child_offset+2].set_range(smem[4 + 1], smem[4 + 2]);
102     children[child_offset+3].set_range(smem[4 + 2], smem[4 + 3]);
103 }

```

**FIGURE 13.14: (Continued)**



```

01 //Some inline vector math functions
02 __forceinline__ __device__ float2 operator+(float2 a, float2 b) {
03     float2 c;
04     c.x = a.x + b.x;    c.y = a.y + b.y;
05     return c;
06 }
07
08 __forceinline__ __device__ float2 operator -(float2 a, float2 b) {
09     float2 c;
10     c.x = a.x - b.x;    c.y = a.y - b.y;
11     return c;
12 }
13
14 __forceinline__ __device__ float2 operator*(float a, float2 b) {
15     float2 c;
16     c.x = a * b.x;    c.y = a * b.y;
17     return c;
18 }
19
20 __forceinline__ __device__ float length(float2 a) {
21     return sqrtf(a.x*a.x + a.y*a.y);
22 }
23
24 //Device function that computes the curvature of a line
25 __device__ float computeCurvature(BezierLine *bLines){
26     int bidx = blockIdx.x;
27     float curvature = length(bLines[bidx].CP[1] - 0.5f*(bLines[bidx].CP[0]
28     + bLines[bidx].CP[2]))/length(bLines[bidx].CP[2]
29     - bLines[bidx].CP[0]);
30     return curvature;
31 }
32
33 void initializeBLines(BezierLine *bLines_h) {
34     //Set initial point to zero (last is last point in the previous segment)
35     float2 last = {0,0};
36     for(int i = 0; i < N_LINES; i++){
37         //Set first point of this line to last point of previous line
38         bLines_h[i].CP[0] = last;
39         for(int j = 1; j < 3; j++) {
40             //Assign random coordinate between 0 and 1
41             bLines_h[i].CP[j].x = (float)rand()/(float)RAND_MAX;
42             //Assign random coordinate between 0 and 1
43             bLines_h[i].CP[j].y = (float)rand()/(float)RAND_MAX;
44         }
45         last = bLines_h[i].CP[2];    //keep the last point of this line
46         //Set number of tessellated vertices to zero
47         bLines_h[i].nVertices = 0;
48     }
49 }

```

**FIGURE A13.8** : Support code for Bezier Curve calculation without dynamic parallelism.

```

01 // A structure of 2D points
02 class Points {
03     float *m_x;
04     float *m_y;
05
06 public:
07     // Constructor
08     __host__ __device__ Points() : m_x(NULL), m_y(NULL) {}
09
10     // Constructor
11     __host__ __device__ Points(float *x, float *y) : m_x(x), m_y(y) {}
12
13     // Get a point
14     __host__ __device__ __forceinline__ float2 get_point(int idx) const {
15         return make_float2(m_x[idx], m_y[idx]);
16     }
17
18     // Set a point
19     __host__ __device__ __forceinline__ void set_point(int idx, const float2 &p) {
20         m_x[idx] = p.x;
21         m_y[idx] = p.y;
22     }
23
24     // Set the pointers
25     __host__ __device__ __forceinline__ void set(float *x, float *y) {
26         m_x = x;
27         m_y = y;
28     }
29 };
30
31 // A 2D bounding box
32 class Bounding_box {
33     // Extreme points of the bounding box
34     float2 m_p_min;
35     float2 m_p_max;
36
37 public:
38     // Constructor. Create a unit box
39     __host__ __device__ Bounding_box(){
40         m_p_min = make_float2(0.0f, 0.0f);
41         m_p_max = make_float2(1.0f, 1.0f);
42     }
43
44     // Compute the center of the bounding-box
45     __host__ __device__ void compute_center(float2 &center) const {
46         center.x = 0.5f * (m_p_min.x + m_p_max.x);
47         center.y = 0.5f * (m_p_min.y + m_p_max.y);
48     }
49
50     // The points of the box
51     __host__ __device__ __forceinline__ const float2 &get_max() const {
52         return m_p_max;
53     }
54
55     __host__ __device__ __forceinline__ const float2 &get_min() const {
56         return m_p_min;
57     }
58
59     // Does a box contain a point
60     __host__ __device__ bool contains(const float2 &p) const {
61         return p.x>=m_p_min.x && p.x<m_p_max.x && p.y>=m_p_min.y && p.y<m_p_max.y;
62     }
63
64     // Define the bounding box
65     __host__ __device__ void set(float min_x, float min_y, float max_x, float max_y){
66         m_p_min.x = min_x;
67         m_p_min.y = min_y;
68         m_p_max.x = max_x;
69         m_p_max.y = max_y;
70     }
71 };

```

**FIGURE A13.13** : Support code for quadtree with dynamic parallelism: definition of points and bounding box.

```

001 // A node of a quadtree
002 class Quadtree_node {
003     // The identifier of the node
004     int m_id;
005     // The bounding box of the tree
006     Bounding_box m_bounding_box;
007     // The range of points
008     int m_begin, m_end;
009
010 public:
011     // Constructor
012     __host__ __device__ Quadtree_node() : m_id(0), m_begin(0), m_end(0) {}
013
014     // The ID of a node at its level
015     __host__ __device__ int id() const {
016         return m_id;
017     }
018
019     // The ID of a node at its level
020     __host__ __device__ void set_id(int new_id) {
021         m_id = new_id;
022     }
023
024     // The bounding box
025     __host__ __device__ __forceinline__ const Bounding_box &bounding_box() const {
026         return m_bounding_box;
027     }
028
029     // Set the bounding box
030     __host__ __device__ __forceinline__ void set_bounding_box(float min_x,
031 float min_y, float max_x, float max_y) {
032         m_bounding_box.set(min_x, min_y, max_x, max_y);
033     }
034
035     // The number of points in the tree
036     __host__ __device__ __forceinline__ int num_points() const {
037         return m_end - m_begin;
038     }
039
040     // The range of points in the tree
041     __host__ __device__ __forceinline__ int points_begin() const {
042         return m_begin;
043     }
044
045     __host__ __device__ __forceinline__ int points_end() const {
046         return m_end;
047     }
048
049     // Define the range for that node
050     __host__ __device__ __forceinline__ void set_range(int begin, int end) {
051         m_begin = begin;
052         m_end = end;
053     }
054 };
055
056 // Algorithm parameters
057 struct Parameters {
058     // Choose the right set of points to use as in/out
059     int point_selector;
060     // The number of nodes at a given level (2^k for level k)
061     int num_nodes_at_this_level;
062     // The recursion depth
063     int depth;
064     // The max value for depth
065     const int max_depth;
066     // The minimum number of points in a node to stop recursion
067     const int min_points_per_node;
068
069     // Constructor set to default values.
070     __host__ __device__ Parameters(int max_depth, int min_points_per_node) :
071         point_selector(0),
072         num_nodes_at_this_level(1),
073         depth(0),
074         max_depth(max_depth),
075         min_points_per_node(min_points_per_node) {}
076

```

**FIGURE A13.14:** Support code for quadtree with dynamic parallelism: definitions and main function.

```

077 // Copy constructor. Changes the values for next iteration
078 __host__ __device__ Parameters(const Parameters &params, bool) :
079     point_selector((params.point_selector+1) % 2),
080     num_nodes_at_this_level(4*params.num_nodes_at_this_level),
081     depth(params.depth+1),
082     max_depth(params.max_depth),
083     min_points_per_node(params.min_points_per_node) {}
084 };
085
086 // Main function
087 void main(int argc, char **argv) {
088
089     // Constants to control the algorithm
090     const int num_points = atoi(argv[0]);
091     const int max_depth = atoi(argv[1]);
092     const int min_points_per_node = atoi(argv[2]);
093
094     // Allocate memory for points
095     thrust::device_vector<float> x_d0(num_points);
096     thrust::device_vector<float> x_d1(num_points);
097     thrust::device_vector<float> y_d0(num_points);
098     thrust::device_vector<float> y_d1(num_points);
099
100     // Generate random points
101     Random_generator rnd;
102     thrust::generate(
103         thrust::make_zip_iterator(thrust::make_tuple(x_d0.begin(), y_d0.begin())),
104         thrust::make_zip_iterator(thrust::make_tuple(x_d0.end(), y_d0.end())),
105         rnd);
106
107     // Host structures to analyze the device ones
108     Points points_init[2];
109     points_init[0].set(thrust::raw_pointer_cast(&x_d0[0]),
110                       thrust::raw_pointer_cast(&y_d0[0]));
111     points_init[1].set(thrust::raw_pointer_cast(&x_d1[0]),
112                       thrust::raw_pointer_cast(&y_d1[0]));
113
114     // Allocate memory to store points
115     Points *points;
116     cudaMalloc((void **) &points, 2*sizeof(Points));
117     cudaMemcpy(points, points_init, 2*sizeof(Points), cudaMemcpyHostToDevice);
118
119     // We could use a close form...
120     int max_nodes = 0;
121
122     for (int i=0, num_nodes_at_level=1; i<max_depth; ++i, num_nodes_at_level*=4)
123         max_nodes += num_nodes_at_level;
124
125     // Allocate memory to store the tree
126     Quadtree_node root;
127     root.set_range(0, num_points);
128     Quadtree_node *nodes;
129     cudaMalloc((void **) &nodes, max_nodes*sizeof(Quadtree_node));
130     cudaMemcpy(nodes, &root, sizeof(Quadtree_node), cudaMemcpyHostToDevice);
131
132     // We set the recursion limit for CDP to max depth
133     cudaDeviceSetLimit(cudaLimitDevRuntimeSyncDepth, max_depth);
134
135     // Build the quadtree
136     Parameters params(max_depth, min_points_per_node);
137     const int NUM_THREADS_PER_BLOCK = 128;
138     const size_t smem_size = 8*sizeof(int);
139     build_quadtree_kernel<<<1, NUM_THREADS_PER_BLOCK, smem_size>>>
140         (nodes, points, params);
141     cudaGetLastError();
142
143     // Free memory
144     cudaFree(nodes);
145     cudaFree(points);
146
147 }

```

**FIGURE A13.14: (Continued)**

```

01 //Some inline vector math functions
02 __forceinline__ __device__ float2 operator+(float2 a, float2 b) {
03     float2 c;
04     c.x = a.x + b.x;    c.y = a.y + b.y;
05     return c;
06 }
07
08 __forceinline__ __device__ float2 operator -(float2 a, float2 b) {
09     float2 c;
10     c.x = a.x - b.x;    c.y = a.y - b.y;
11     return c;
12 }
13
14 __forceinline__ __device__ float2 operator*(float a, float2 b) {
15     float2 c;
16     c.x = a * b.x;    c.y = a * b.y;
17     return c;
18 }
19
20 __forceinline__ __device__ float length(float2 a) {
21     return sqrtf(a.x*a.x + a.y*a.y);
22 }
23
24 //Device function that computes the curvature of a line
25 __device__ float computeCurvature(BezierLine *bLines){
26     int bidx = blockIdx.x;
27     float curvature = length(bLines[bidx].CP[1] - 0.5f*(bLines[bidx].CP[0]
28         + bLines[bidx].CP[2]))/length(bLines[bidx].CP[2]
29         - bLines[bidx].CP[0]);
30     return curvature;
31 }
32
33 void initializeBLines(BezierLine *bLines_h) {
34     //Set initial point to zero (last is last point in the previous segment)
35     float2 last = (0,0);
36     for(int i = 0; i < N_LINES; i++){
37         //Set first point of this line to last point of previous line
38         bLines_h[i].CP[0] = last;
39         for(int j = 1; j < 3; j++) {
40             //Assign random coordinate between 0 and 1
41             bLines_h[i].CP[j].x = (float)rand()/(float)RAND_MAX;
42             //Assign random coordinate between 0 and 1
43             bLines_h[i].CP[j].y = (float)rand()/(float)RAND_MAX;
44         }
45         last = bLines_h[i].CP[2]; //Keep the last point of this line
46         //Set number of tessellated vertices to zero
47         bLines_h[i].nVertices = 0;
48     }
49 }
50
51 // A structure of 2D points
52 class Points {
53     float *m_x;
54     float *m_y;
55
56 public:
57     // Constructor
58     __host__ __device__ Points() : m_x(NULL), m_y(NULL) {}
59
60     // Constructor
61     __host__ __device__ Points(float *x, float *y) : m_x(x), m_y(y) {}
62
63     // Get a point
64     __host__ __device__ __forceinline__ float2 get_point(int idx) const {
65         return make_float2(m_x[idx], m_y[idx]);
66     }
67
68     // Set a point
69     __host__ __device__ __forceinline__ void set_point(int idx, const float2 &p) {
70         m_x[idx] = p.x;
71         m_y[idx] = p.y;
72     }
73 }

```

```

24     // Set the pointers
25     __host__ __device__ __forceinline__ void set(float *x, float *y) {
26         m_x = x;
27         m_y = y;
28     }
29 };
30
31 // A 2D bounding box
32 class Bounding_box {
33     // Extreme points of the bounding box
34     float2 m_p_min;
35     float2 m_p_max;
36
37 public:
38     // Constructor. Create a unit box
39     __host__ __device__ Bounding_box(){
40         m_p_min = make_float2(0.0f, 0.0f);
41         m_p_max = make_float2(1.0f, 1.0f);
42     }
43
44     // Compute the center of the bounding-box
45     __host__ __device__ void compute_center(float2 &center) const {
46         center.x = 0.5f * (m_p_min.x + m_p_max.x);
47         center.y = 0.5f * (m_p_min.y + m_p_max.y);
48     }
49
50     // The points of the box
51     __host__ __device__ __forceinline__ const float2 &get_max() const {
52         return m_p_max;
53     }
54
55     __host__ __device__ __forceinline__ const float2 &get_min() const {
56         return m_p_min;
57     }
58
59     // Does a box contain a point
60     __host__ __device__ bool contains(const float2 &p) const {
61         return p.x>=m_p_min.x && p.x<m_p_max.x && p.y>=m_p_min.y && p.y<m_p_max.y;
62     }
63
64     // Define the bounding box
65     __host__ __device__ void set(float min_x, float min_y, float max_x, float max_y){
66         m_p_min.x = min_x;
67         m_p_min.y = min_y;
68         m_p_max.x = max_x;
69         m_p_max.y = max_y;
70     }
71 };

```