

# Brook: A Streaming Programming Language

I. Buck

8 October 2001

## 1 Motivation

### 1.1 Why design our own language?

- Modern processors can be limited by communication rather than computation. However, communication or dependency information can often be difficult to determine at compile time and therefore limit parallelism. Communication and dependency information should be explicit within Brook to alleviate this problem.
- Scientific computing can benefit greatly from parallel computing. However, porting large applications can be difficult if not impossible for unconventional languages. Brook should cater to the scientific computing community and be simple to port from existing C applications.
- Straightforward to parallelize the computation. The programmer should have a clear understanding of what portions of the calculations are going to be parallelized by the system. Both data and task parallelism should be fairly intuitive.
- Retargetable. Mappable to a variety of hardware including existing stream machines (a.k.a Imagine), SMP systems, and future graphics hardware. It should also be clear to the programmer how their code may be mapped to hardware.

## 2 Streams, Stream Functions, and Kernels Functions

### 2.1 What is a Stream?

Brook introduces a new datatype called a **Stream**. A stream is simply a collection of data which can be operated on in parallel. Each stream element consists of a record of values. The layout of these records within the hardware memory is hidden from the programmer to accommodate for the variety of hardware implementations.

```
vec4f a = {...};
vec4f b;
mat4f M = {...};

b = M * a;
```

Figure 1: Transforming vector a by matrix M and storing the result in b.

## 2.2 Stream Functions

The elements of a stream can only be accessed via a stream function. A stream function is a function which is applied to the elements of a stream. A function can take multiple stream inputs and have multiple stream outputs.

## 2.3 Kernel Functions

A kernel function is a type of stream function which operates in parallel over the elements of a stream. The operations allowed in kernel functions are limited in order to insure parallel operation.

# 3 Programming Model

## 3.1 Language Basics

The syntax of Brook is designed to be similar to C with a few additions. A native vector library is provided which offers mathematical types and operations common to scientific computing. These types include: *Vectors*: `vec2f` (2 component, IEEE standard floating point), `vec3f`, `vec4f`, `vec2d` (double precision), `vec3d`, `vec4d`, `vec2i` (32 bit integer), `vec3i`, `vec4i`; *Matrices*: `mat2f` (2x2 matrix), `mat3f`, `mat4f`, `mat2d`, `mat3d`, `mat4d`, `mat2i`, `mat3i`, `mat4i`.

Standard mathematical operations such as computing the dot product or performing a matrix transform are expressed in a C-like style as shown in figure 3.1. The compiler converts these operations to the most efficient hardware implementation.

## 3.2 Stream Declaration

A stream consists of a collection of elements which can be declared similar to C structs. Each element can contain any native C type (float, int, double) or vector type (`vec3f`, `mat4d`). Arrays are allowed although no pointers can exist within a stream declaration. An example stream definition and declaration is shown in figure 3.2.

Streams can be declared as a simple set such as the `vtx` example. The length of the stream is unspecified, allowing it to grow or shrink as needed. Streams

```

\\ Stream definitions
stream Vertex {
    vec4f pos;
    vec4f color;
    vec3f normal;
    vec2f texture;
}

stream Color {
    vec3f c;
}

\\ Stream declaration
Vertex vtx;
Color framebuffer[1024][1024];

```

Figure 2: Defining and declaring a stream. Defining streams is similar to defining a C struct.

can also be declared as arrays, giving them a fixed size as in the *framebuffer* example.

The layout of streams in memory is not exposed to the programmer. The only way to access stream elements is through stream functions.

### 3.3 Stream Functions

Stream functions are functions which process elements of a stream. These functions are defined much like a C function. The inputs of a stream function are elements from a stream as well as any constant values. Below is a simple function which prints the position field of a Vertex stream.

Invoking the stream function *PrintVertexPos* causes the function to be called on each element within the stream *v*. The element passed into the stream function is both readable as well as writable. General stream functions provide a convenient method for unrestricted access to elements within a stream since they operate similarly to C functions. They are capable of conditionals, looping, local variables, accessing global memory, perform function calls, etc. However, they are serial operations and are executed on a single processor or host CPU. For parallel performance, the programmer should use a special type of stream function called a **kernel function**.

### 3.4 Kernel Functions

General stream functions such as *PrintVertexPos* are not parallelized by the compiler nor do they use special streaming hardware within the host system. In

```

void
PrintVertexPos (Vertex vtx) {
    printf ('pos: %f %f %f\n', vtx.pos[0], vtx.pos[1], vtx.pos[2]);
}

void main (void) {
    Vertex v;
    \\ Do some calculations on v
    ....
    \\ Print the Vertex stream
    PrintVertexPos (v);
}

```

Figure 3: A simple stream function.

```

kernel void
vtransform (Vertex vtx, Vertex out tvtx, mat4f matrix) {
    tvtx.v = matrix * vtx.v
}

```

Figure 4: A simple kernel function.

order for stream functions to take advantage of parallel hardware the programmer should use a kernel function.

A kernel functions are a subset of stream functions which allow the use parallel hardware units to operate on a stream in parallel. In order to allow the kernel function to be parallelized, restrictions are placed on the types of operations it is able to perform.

Figure 3.4 shows is an example of a kernel function which applies a matrix transform to each element of vertex stream.

The **kernel** keyword signals to the compiler that this is a kernel function which should be run in parallel on the hardware. Kernel functions, in general, can perform normal C-like operations: Local variables, conditionals, and loops. However certain restrictions are enforced when defining kernels:

1. Global variables are not visible within kernels. Only the arguments of kernel are accessible.
2. Function calls are permitted to other kernel functions only. This prevents the kernels from calling system functions, memory allocation functions, etc.
3. Stream elements can only be accessed read only or write only. Read modify write is not allowed. By default, arguments are considered to be read only, unless the **out** or **outm** keyword is used. (More on these keywords below.)

These restrictions are checked by the compiler if the *kernel* keyword is specified at beginning of the function declaration. One of the key restrictions for kernel function is the read or write only arguments. This allows the compiler to easily detect communication between kernels and build a basic flow graph for the computation. The compiler is also able to schedule the execution of stream functions, both kernel and general stream functions.

### 3.5 1-N Kernels

In the *vtransform* example, there is a one to one mapping from input elements to output elements. One vertex is passed in from *vtx* and one vertex is outputted to *vtvx*. This is the defined behavior for the **out** keyword.

Brook also supports multiple outputs to a stream. Figure 3.5 is a kernel for doing x-major line rasterization where two endpoints of a line are inputted and a variable number of fragments are outputted.

In this example, the *lineraster* program can output multiple fragments which are determined by the line length. The fragment stream is passed in via the argument list with the **outm** keyword. This tells the compiler that the kernel may output a variable number of elements to stream *f*, including no elements at all. To output to the stream, the kernel uses the **push** function which outputs the current value of *f* to the stream.

To summarize, there are two different types of keywords for outputs.

1. **out**: One output element is generated per input element. **push** operator is not required.
2. **outm**: Zero to n elements generated for an input element. Each element must be explicitly outputted using the **push** function.

## 4 Array Accesses

In the previous examples, streams have been declared as simple sets. For computations which require convolution operations the programmer may want to query the neighboring values within a stream.

To allow for this, Brook supports neighborhood addressing. Figure 4 demonstrates a 3x3 convolution kernel which performs a simple blur at given a rate. For each element, it computes a new value based on the weighted sum of itself with the value above, below, left, and right.

In this example, the argument list to the kernel function includes the relative addressing information. *Convolve* takes a stream input in the form of the array. *x* and *y* are **offset variables** which contain the current offset of *grid\_in* which is being processed by the kernel function. Following the offset variables is the **range information**. In this example, the kernel declares that it access one element left and right of *x* ( $[x:-1,1]$ ) and one element above and below of *y* ( $[y:-1,1]$ ). The compiler uses this information to know what parts of *grid\_in* must be defined before calling the function. The range can be specified for

```

stream Line {
    vec2i a;
    vec2i b;
}

stream Fragment {
    vec2i pos;
}

kernel void
linerafter (Line l, Fragment outm f) {
    int dx = abs (l.a.x - l.b.x), dy = abs (l.a.y - l.b.y);
    int p = 2 * dy - dx;
    int twoDy = 2 * dy, twoDyDx = 2 * (dy - dx);
    int xEnd;
    vec2i v;

    /* Determine which point to use as start, which as end */
    if (l.a.x > l.b.x) {
        v = l.b;
        xEnd = l.a.x;
    } else {
        v = l.a;
        xEnd = l.b.x
    }

    /* Write the first fragment */
    f.pos = v;
    push(f);

    /* Bresenham loop */
    while (v.x < xEnd) {
        v.x++;
        if (p < 0) p += twoDy;
        else { v.y++; p += twoDyDx; }

        f.pos = v
        push(f);
    }
}

```

Figure 5: A xmajor line rasterization kernel which demonstrates the multiple element output from a kernel using the push operator.

```

stream Flow {
    float t;
}

kernel void
SetValues (Flow out grid_out[x][y]) {
    grid_out.t = x + y;
}

kernel void
Convolve (Flow grid_in[x:-1,1][y:-1,1], Flow out grid_out,
          float rate) {

    float k = 1.0f - (rate*4.0f);
    float a, b, c, d;
    float t = grid_in.t;

    /* Check the boundry condition */
    a = (x==0)      ? t : grid_in[-1][0].t;
    b = (x==XMAX) ? t : grid_in[+1][0].t;
    c = (y==0)      ? t : grid_in[0][-1].t;
    d = (y==YMAX) ? t : grid_in[0][+1].t;

    grid_out.t =
        t * k +
        a * rate + b * rate +
        c * rate + d * rate;
}

void main (void) {
    Flow grid[1024][1024];

    /* initialize grid */
    SetValues(grid);

    /* Perform the convolution */
    Convolve (grid, grid, 0.1f)
}

```

Figure 6: A relative addressing example.

```

kernel void
Convolve (Flow grid_anyx_in[x:][y:-1,1], Flow grid_out, float rate) {
....

```

Figure 7: Example of absolute range array. Here the x component of grid\_anyx\_in is an absolute address.

both outputs and inputs if needed. Accessing the elements of the array is done through relative offsets from the offset variables.

In some cases, a kernel may only access the  $x,y$  value of an array and not require any other values. This is similar to the range  $[x:0,0]$ . In this case, the range information does not need to be specified, as shown in the *SetValues* kernel function. The  $:0,0$  is not required and the zero range is assumed.

Finally, a kernel may require absolute addressing within an array. This is the case if the offset into the array is computed by a function and is not known at compile time. The range is therefore the entire array. This is specified by not indicating a range after the  $:$  character, e.g.  $[x:]$ . Figure 4 declares function argument that can be accessed anywhere in the x direction but only one element above or below in the y direction. When accessing *grid\_anyx\_in*, the x component is absolute instead of relative.

The programmer should try to specify the minimum range necessary for the function to operate. If all the offset variables are specified as absolute then the full array must be defined in order for the function to be executed. This can limit the task parallelism of the kernel.

## 5 Reduction Operations

Stream function includes a return value which can be used to gather information about the stream computation. Often with iterative scientific methods a particular computation is repeated until a system stabilizes. Figure 5 shows a kernel which returns the difference of a value amongst its neighbors. The main loop repeats the convolution kernel until the difference is below a certain threshold.

The **max** operator returns the maximum value from all the Divergence kernel operations. (MORE TEXT HERE)

## 6 Stream Management

(MORE TEXT HERE)

## 7 Special Streams

There are some special stream types defined in Brook .



```

kernel float
Divergence (Flow grid_in[x:-1,1][y:-1,1]) {

    float a, b, c, d;
    float t = grid_in.t;
    float div;

    a = (x==0) ? t : grid_in[-1][0].t;
    b = (x==XMAX)? t : grid_in[+1][0].t;
    c = (y==0) ? t : grid_in[0][-1].t;
    d = (y==YMAX)? t : grid_in[0][+1].t;

    return abs(t-a) + abs(t-b) +
           abs(t-c) + abs(t-d);
}

void main (void) {
Flow grid[XMAX][YMAX];
const float rate = 0.1;
float diverge;
do {
    Convolve (grid, grid, rate);
    diverge = max Divergence (grid);
} while (diverge > 5.0f);
}

```

Figure 8: Example of a reduction operation. The kernel function returns a floating point number. The **max** operator returns the maximum value returned from the Divergence function.

1. **FileStream**: A file stream which can be read or written with standard `stdio.h` functions like `fscanf`, `fprintf`, `fread`. Equivalent to a `FILE` pointer in C. These streams cannot be passed to kernel functions.
2. **(MORE TEXT HERE)**