

用 GPU 加速运行生命游戏

2021 年 1 月 23 日

1 引言

在一个长宽给定的二维数组 (array) 中，将数组的每个元素视作一个「细胞」，若一个元素的值为 1，代表这个细胞「存活」，否则代表这个细胞不存活，在一次迭代中，对该数组的每一个元素，考虑与它最相邻的 8 个元素，将这 8 个元素的值求和，代表这个细胞它存活的邻居的数量，如果存活的邻居数量少于 2，那么这个细胞死去（通过对它置 0）；如果存活的邻居数量刚好是 2 或者 3，那么这个细胞继续存活；如果存活的邻居的数量超过 3，那这个细胞因为「养分不足」而死去；如果一个已经死去的细胞周围刚好有 3 个存活着的细胞，那么它重生（对它置 1）。这样的计算规则叫做「生命游戏」(Conway's Game of Life)，他是相同维数的数组到数组的对应法则，并且每次计算只考虑每一个元素的「邻居」，这样的对应法则统称为元胞自动机 (cellular automaton)。在每一轮迭代更新中，所有元素都是被同步更新的，所以元胞自动机天然适合在 GPU 上实现。

2 实现

首先我们确定故事的主角，也就是那个二维数组的初始状态，这个可以随意设置：

```
1 import cupy as cp
2
3 initial_state = cp.array([
4     [0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0],
5     [0,0,0,0,0,1,1,0,0,0,0,0,0,1,1,0,0,0,0,0],
6     [0,0,0,0,1,0,1,0,0,0,0,0,0,1,0,1,0,0,0,0],
7     [0,0,0,0,1,0,0,0,0,0,0,0,0,0,0,1,0,0,0,0],
8     [0,1,1,0,1,0,0,0,0,0,0,0,0,0,0,1,0,1,1,0],
9     [0,1,1,0,1,0,1,0,0,1,1,0,0,1,0,1,0,1,1,0],
10    [0,0,0,0,1,0,1,0,1,0,0,1,0,1,0,1,0,0,0,0],
11    [0,0,0,0,1,0,1,0,1,0,0,1,0,1,0,1,0,0,0,0],
12    [0,1,1,0,1,0,1,0,0,1,1,0,0,1,0,1,0,1,1,0],
13    [0,1,1,0,1,0,0,0,0,0,0,0,0,0,0,1,0,1,1,0],
14    [0,0,0,0,1,0,0,0,0,0,0,0,0,0,0,1,0,0,0,0],
15    [0,0,0,0,1,0,1,0,0,0,0,0,0,0,1,0,1,0,0,0],
16    [0,0,0,0,0,1,1,0,0,0,0,0,0,1,1,0,0,0,0,0],
17    [0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0]
18 ])
```

那么初始状态确定了，因为生命游戏是不改变数组的长宽的，所以数组的长宽也就确定了，顺便确定好要演化多少代：

```
1 n_genes = 10
```

```

2 n_rows = initial_state.shape[0]
3 n_cols = initial_state.shape[1]

```

我们计划为每一个细胞分配一个 thread 去更新，所以要提前算好资源的分配量：

```

1 from math import ceil
2
3 BLOCK_SIZE = 16
4 block_dim = (BLOCK_SIZE, BLOCK_SIZE,)
5 grid_dim = (ceil(n_cols/block_dim[0]), ceil(n_rows/block_dim[1]),)

```

然后，我们生成一个 `n_genes` 层，`n_rows` 行，`n_cols` 列的高维数组，每一代存储在每一层里面，并且将第一层置为初始状态：

```

1 data = cp.zeros(shape=(n_genes+1, n_rows, n_cols,), dtype=cp.uint32)
2 data[0, :, :] = initial_state

```

站在一个 thread 的视角上，按照生命游戏的规则去编写每个 thread 的工作剧本：

```

1 evolve = cp.RawKernel(
2     r'''
3     #include <cooperative_groups.h>
4
5     using namespace cooperative_groups;
6
7     extern "C" __global__
8     void evolve(unsigned int *data, int n_rows, int n_cols, int max_generation)
9     {
10         if (sizeof(unsigned int) != 4)
11         {
12             return;
13         }
14
15         int row = (unsigned int) (blockIdx.y * blockDim.y + threadIdx.y);
16         int col = (unsigned int) (blockIdx.x * blockDim.x + threadIdx.x);
17
18         if ((row >= n_rows) || (col >= n_cols))
19         {
20             return;
21         }
22
23         for (int gen = 0; gen < max_generation; ++gen)
24         {
25             int l_row = row;
26             int l_col = (col - 1) % n_cols;
27
28             int r_row = row;
29             int r_col = (col + 1) % n_cols;
30
31             int u_row = (row - 1) % n_rows;
32             int u_col = col;
33
34             int d_row = (row + 1) % n_rows;
35             int d_col = col;

```

```

36
37     int ul_row = (row - 1) % n_rows;
38     int ul_col = (col - 1) % n_cols;
39
40     int ur_row = (row - 1) % n_rows;
41     int ur_col = (col + 1) % n_cols;
42
43     int dr_row = (row + 1) % n_rows;
44     int dr_col = (col + 1) % n_cols;
45
46     int dl_row = (row + 1) % n_rows;
47     int dl_col = (col - 1) % n_cols;
48
49     unsigned n_neighbors = 0;
50
51     n_neighbors += data[gen * n_rows * n_cols + u_row * n_cols + u_col];
52     n_neighbors += data[gen * n_rows * n_cols + d_row * n_cols + d_col];
53     n_neighbors += data[gen * n_rows * n_cols + l_row * n_cols + l_col];
54     n_neighbors += data[gen * n_rows * n_cols + r_row * n_cols + r_col];
55
56     n_neighbors += data[gen * n_rows * n_cols + ur_row * n_cols + ur_col];
57     n_neighbors += data[gen * n_rows * n_cols + ul_row * n_cols + ul_col];
58     n_neighbors += data[gen * n_rows * n_cols + dr_row * n_cols + dr_col];
59     n_neighbors += data[gen * n_rows * n_cols + dl_row * n_cols + dl_col];
60
61     unsigned int this_index = gen * n_rows * n_cols + row * n_cols + col;
62     unsigned int next_index = this_index + n_rows * n_cols;
63
64     if (n_neighbors < 2 && data[this_index] == 1)
65     {
66         data[next_index] = 0;
67     }
68     else if ((n_neighbors == 2 || n_neighbors == 3) && data[this_index] == 1)
69     {
70         data[next_index] = 1;
71     }
72     else if (n_neighbors > 3 && data[this_index] == 1)
73     {
74         data[next_index] = 0;
75     }
76     else if (n_neighbors == 3 && data[this_index] == 0)
77     {
78         data[next_index] = 1;
79     }
80     else
81     {
82         data[next_index] = data[this_index];
83     }
84
85     grid_group grid = this_grid();
86     grid.sync();

```

```

87     }
88 }
89 '''
90 'evolve',
91 backend='nvcc',
92 enable_cooperative_groups=True
93 )

```

用到 `grid.sync()` 的原因是每个 `thread` 在更新每个细胞时，都要考虑这个细胞周围细胞的状态，所以必须得等所有细胞都完工才能进入下一代。一切都准备好后就可以开始执行了：

```

evolve(grid_dim, block_dim, (data, n_rows, n_cols, n_genes))

```

通过将数组打印出来我们可以看到演化过程：

```

1 import matplotlib.pyplot as plt
2
3 fig, axs = plt.subplots(2, 3)
4
5 for i in range(2):
6     for j in range(3):
7         k = i * 3 + j
8         axs[i, j].matshow(data[k, :, :].get())
9         axs[i, j].axis('off')
10        axs[i, j].set_title(f"Gene: {k}")
11
12 plt.savefig('game-of-life.pdf', bbox_inches='tight')

```

输出见图 2-1

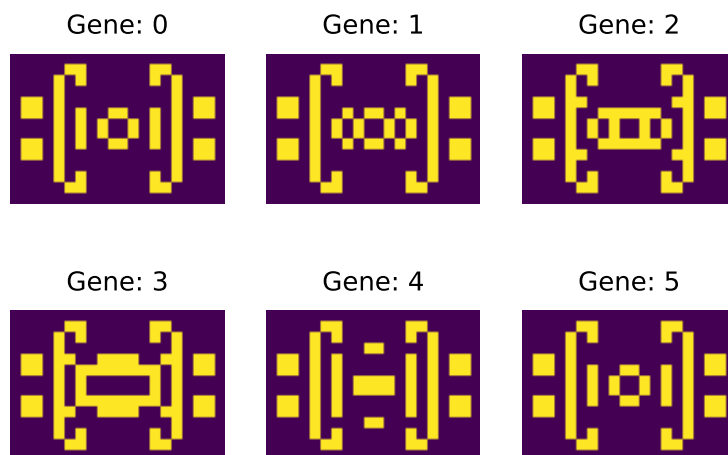


图 2-1: 元胞自动机演化过程

会看到第 6 代也就是 `data[5, :, :]` 和第 1 代 `data[0, :, :]` 是相同的，这说明如果选择某些特殊的初始状态，元胞自动机将会展现出周期性。