

## Project 3 Experiments

The following queries were used to test our current implementation:

Query1: SELECT \* FROM Candy WHERE Candy.C < 505;

Query2: SELECT \* FROM Candy WHERE Candy.C < 750;

Query3: SELECT \* FROM Candy, Pudding WHERE Candy.C < 505 AND  
Pudding.E < 5;

Data has the following properties and were generated as follows:

Relation 1: Candy

Schema: Candy A B C D

Number of tuples: 7500

Attributes were all chosen uniformly at random in the following ranges:  
0 to 10, 0 to 100, 500 to 1000, and 500 to 1000, respectively.

Relation 2: Pudding

Schema: Pudding E F G H

Number of tuples: 5000

Attributes were all chosen uniformly at random in the following ranges:  
0 to 100, 0 to 100, 50 to 100, and 500 to 1000, respectively.

The indexes used for each query were all on Candy.C.

The join algorithm used was BNLJ with a buffer size of 4.

Bar graphs of running times are on the following page. A log base y-scale graph is also included.

