

Scalable Parallelization of Skyline Computation for Multi-core Processors

堀 川 隆

概要

- 趣旨
 - Skyline Computationを並列実行するアルゴリズムを提案, 評価
- 提案アルゴリズム
 - Q-Flow : データ分割は単純な方法(データをsortして分割)で行い, 各々を並列に処理 (アルゴリズムの骨子説明用?)
 - Hybrid : pivot となるデータを選び, それに基づいてデータ分割を行った上で, Q-Flowと同様の手法で並列処理
- ポイント
 - 独立して実行できるDominance test (2点を比較し, 片方が優勢かどうかを判定)を並列に実行
 - Dominance testの実行回数を削減

アルゴリズム (Q-Flow)

Algorithm 1 Q-Flow ($\mathcal{P}, \alpha \rightarrow \text{SKY}(\mathcal{P})$)


```

1:  $S \leftarrow \emptyset$ .
2: Prefilter, sort, and/or partition  $\mathcal{P}$ .           ▷ Initialization
3: while  $\mathcal{P} \neq \emptyset$  do
4:    $Q \leftarrow$  next  $\alpha$  points of  $\mathcal{P}$ .
5:    $\mathcal{P} \leftarrow \mathcal{P} \setminus Q$ .
6:   for  $i \in [0, |Q|)$  do                         ▷ Parallel Phase I
7:     if  $\exists q \in S : q \prec Q[i]$  then
8:       Mark  $Q[i]$  as pruned.
9:   Remove pruned  $Q[i]$  from  $Q$ .                   ▷ Compression
10:  for  $i \in [0, |Q|)$  do                           ▷ Parallel phase II
11:    if  $\exists j \in [0, i) : Q[j] \prec Q[i]$  then
12:      Mark  $Q[i]$  as pruned.
13:  Remove pruned  $Q[i]$  from  $Q$ .                   ▷ Compression
14:  Append  $Q$  to  $S$ .                               ▷ Update structure
15: return  $S$ .
  
```


← データをQ, P, ...に分割

← 各点を既判明のSkylineとDT

← 残った点をブロック内の点とDT



(a) Q-Flow Phase I: comparing to skyline points.



(b) Q-Flow Phase II: comparing to other survivors.

(図は総て論文から引用, 以下同様)

アルゴリズム (Hybrid)

Algorithm 1 Q-Flow ($\mathcal{P}, \alpha \rightarrow \text{SKY}(\mathcal{P})$)

```

1:  $S \leftarrow \emptyset$ .
2: Prefilter, sort, and/or partition  $\mathcal{P}$ .
3: while  $\mathcal{P} \neq \emptyset$  do
4:    $Q \leftarrow$  next  $\alpha$  points of  $\mathcal{P}$ .
5:    $\mathcal{P} \leftarrow \mathcal{P} \setminus Q$ .
6:   for  $i \in [0, |Q|)$  do
7:     if  $\exists q \in S : q \prec Q[i]$  then
8:       Mark  $Q[i]$  as pruned.
9:   Remove pruned  $Q[i]$  from  $Q$ .
10:  for  $i \in [0, |Q|)$  do
11:    if  $\exists j \in [0, i) : Q[j] \prec Q[i]$  then
12:      Mark  $Q[i]$  as pruned.
13:  Remove pruned  $Q[i]$  from  $Q$ .
14:  Append  $Q$  to  $S$ .
15: return  $S$ .
```

▷ Initialization

▷ Parallel Phase I

▷ Compression

▷ Parallel phase II

▷ Compression

▷ Update structure

Algorithm 2 updateS&M($S, \mathcal{M}(S), Q$)

```

1: Append  $Q$  to  $S$ .
2: Pop sentinel off of  $\mathcal{M}(S)$ .
3:  $(m, i) \leftarrow \text{top}(\mathcal{M}(S))$ 
4: for  $j \in [0, |Q|)$  do
5:   if  $Q[j].m = m$  then
6:      $Q[j].m \leftarrow \text{part}(Q[j], S[i])$ .
7:   else
8:      $(m, i) \leftarrow (Q[j].m, |S| + j)$ .
9:     Push  $(m, i)$  onto  $\mathcal{M}(S)$ .
10: Push sentinel ( $2^d, |S| + |Q|$ ) onto  $\mathcal{M}(S)$ .
```

Algorithm 3 compareToSky ($S, \mathcal{M}(S), q \rightarrow \{\text{true}, \text{false}\}$)

```

1: for  $i \in [0, |\mathcal{M}(S)|)$  do
2:    $m \leftarrow \mathcal{M}(S)[i].m$ .
3:   if  $m$  is not incomparable to  $q.m$  then
4:      $s \leftarrow \mathcal{M}(S)[i].s; t \leftarrow \mathcal{M}(S)[i+1].s$ 
5:      $m' \leftarrow \text{part}(q, S[s])$ .
6:     if  $m' = 2^d - 1 \wedge q \neq S[s]$  then return true.
7:     for  $j \in (s, t)$  do
8:       if  $S[j].m$  is not incomparable to  $m'$  then
9:         if  $S[j] \prec q$  then return true.
10: return false.
```

Algorithm 4 compareToPeers ($Q, me \rightarrow \{\text{true}, \text{false}\}$)

```

1:  $i \leftarrow 0$ .
2: while  $Q[i].|m| < Q[me].|m|$  do
3:   if  $Q[i]$  is not pruned then
4:     if  $Q[i].m$  is not incomparable to  $Q[me].m$  then
5:       if  $Q[i] \prec Q[me]$  then return true.
6:    $i \leftarrow i + 1$ .
7: while  $Q[i].m < Q[me].m$  do
8:    $i \leftarrow i + 1$ .
9: while  $i < me$  do
10:  if  $Q[i] \prec Q[me]$  then return true.
11: return false.
```

評価

- 比較対象
 - Pskyline : the state-of-the-art multicore algorithm
 - BSkyTree : the state-of-the-art sequential algorithm
 - PBSkyTree : BSkyTreeを筆者らが並列化したもの
- 環境
 - マシン
 - CPU : Xeon E5-2670(8 cores) x 2, Hyper Threading off ⇒ 16 cores
 - メモリ: 64 GB
 - データセット
 - Standard skyline data generator により, correlated, independent, anticorrelated の3種類を生成
 - Dimension: 4 ~ 16 (default 12)
 - Cardinality: 500K ~ 8M (default 1M)
 - 実データ(参考用) : NBA, House (比較的小規模) と Weather(大規模)

結果 (実行時間)

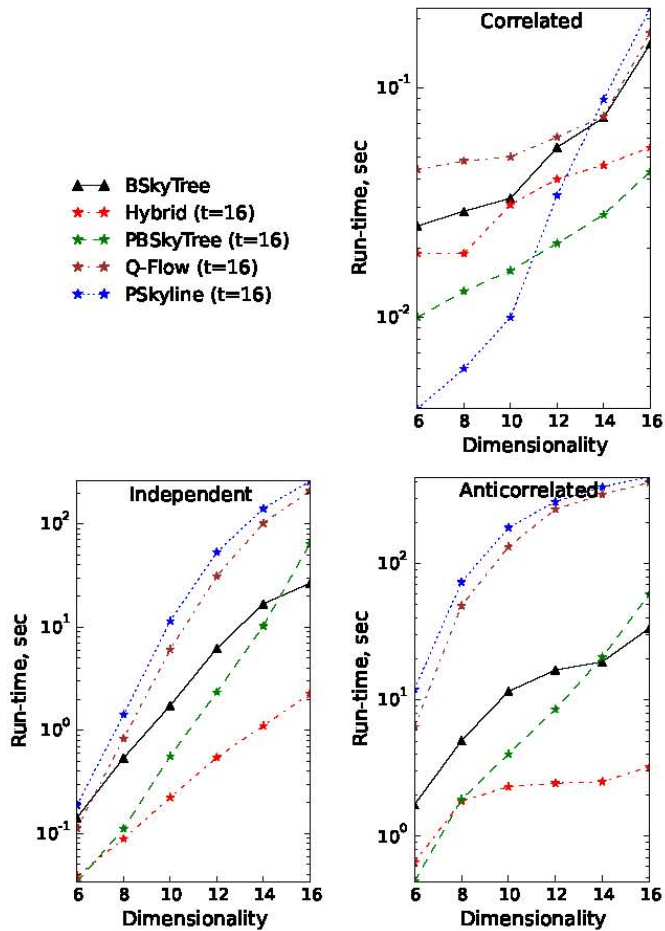


Fig. 5: State-of-the-art performance w.r.t. d ($n = 1M$).

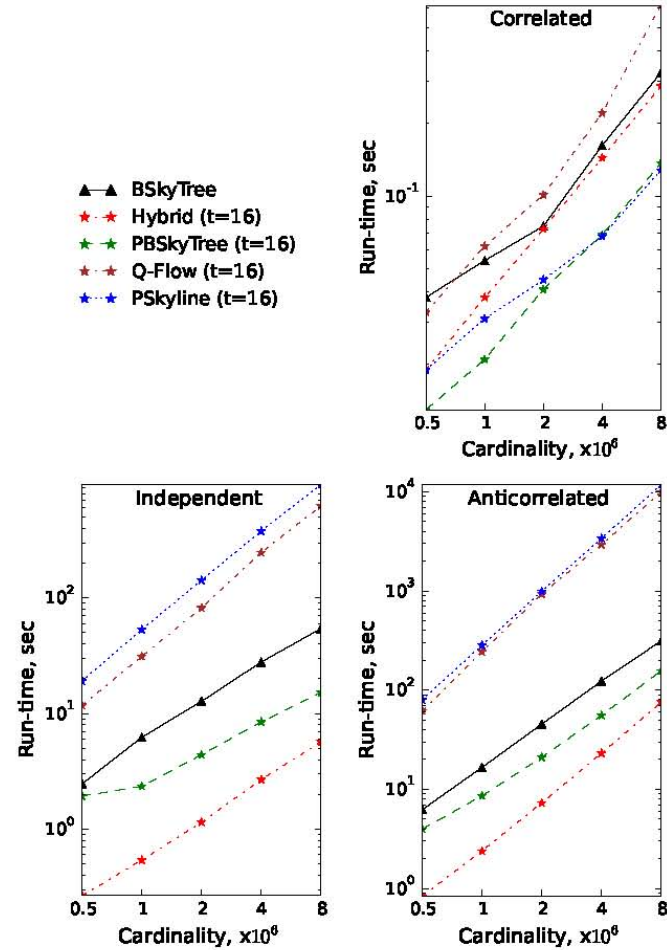


Fig. 6: State-of-the-art performance w.r.t. n ($d = 12$).

処理時間の増え方は、データの種類によって異なる

結果（実行時間のbreak down）




Fig. 7: Effect of α in Q-Flow with varied α ($n = 1M$, $d = 12$)

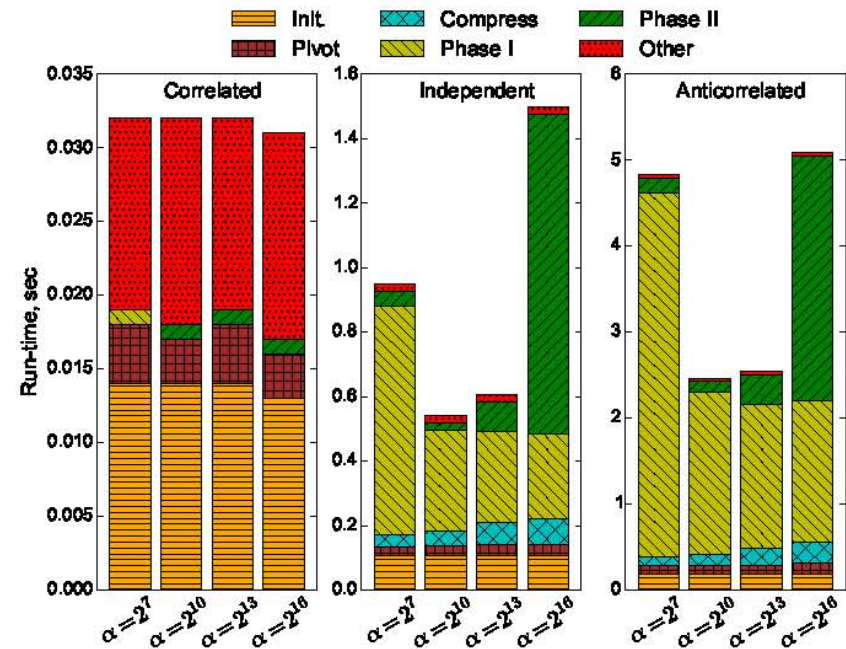


Fig. 8: Effect of α on Hybrid with varied α ($n = 1M$, $d = 12$)

- Correlatedでは初期化時間の割合が大きい。他のデータでは、初期化時間の増加を補って余りある並列化効果が出ている。

まとめ

- Skyline Computationの並列アルゴリズムを考案
 - Q-Flow: 最先端の並列アルゴリズム (Pskyline) を上回る
 - Hybrid: partitioningを行う版, これまでの内, 最も効率的とされているアルゴリズム (BSkyTree), および, その並列化版 (PBSkyTree) を上回る