A General Cache Framework for Efficient Generation of Timing Critical Paths

Kuan-Ming Lai, Tsung-Wei Huang and Tsung-Yi Ho

Department of Computer Science, National Tsing Hua University, Hsinchu, Taiwan

Department of Electrical and Computer Engineering, University of Illinois at Urbana-Champaign, IL, USA

r365460@gapp.nthu.edu.tw, twh760812@gmail.com, tyho@cs.nthu.edu.tw

ABSTRACT

The recent TAU 2018 contest was seeking novel idea for efficient generation of timing reports. When the timing graph is updated, users query different forms of timing reports that happen subsequently and sequentially. This process is computationally expensive and inherently complex. Therefore, we introduce in this paper a general cache framework for efficient generation of timing critical paths. Our framework efficiently supports (1) a cache scheme to minimize duplicate calculation, (2) graph contraction to reduce the search space, and (3) multi-threading. We evaluated our framework on the TAU 2018 contest benchmarks and demonstrated promising performance over the top performer.

KEYWORDS

Static Timing Analysis, Path-based Timing Analysis, Cache

ACM Reference Format:

Kuan-Ming Lai, Tsung-Wei Huang and Tsung-Yi Ho. 2019. A General Cache Framework for Efficient Generation of Timing Critical Paths. In *The 56th Annual Design Automation Conference 2019 (DAC '19)*, June 2–6, 2019, Las Vegas, NV, USA. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3316781.3317744>

1 INTRODUCTION

Timing-driven optimization is an important step in many design flows such as logic synthesis, placement, routing, and physical synthesis [1]. To achieve timing closure, the chip designers iteratively optimize the timing of their designs until all timing constraints are satisfied. Optimization tools often call a timer in their inner loop to report timing critical paths to evaluate a strategy, for example, changing the size of a gate, adding a buffer, or removing a gate. The timer must quickly and accurately update the timing to ensure slack integrity and timing closure for reasonable turnaround time and performance.

One of the core routines in implementing an efficient timing analysis tool is the generation of timing critical paths. In many timers, this often refers to the command or the application programming interface (API) method “report_timing”. The algorithm behind this command is critical to the performance of using it in

Figure 1: An example of timing critical path report queries.

a timing-driven optimization loop. A typical user flow can have more than 10K commands that request timing reports for a specific subset of paths [2]. Depending on applications, there are many aspects and requirements to this command. A common feature is to generate the top- k timing critical path on a specific cone of logic. This has been addressed in the recent TAU 2018 Timing Analysis Contest [2].

Figure 1 shows an example of two queries to report timing critical paths in one incremental timing iteration. Query 1 requests the critical path to go through pin B and pin J (modeled as a graph node). Query 2 requests the critical path to go through pin A and pin J. An intuitive solution is to implement a standalone algorithm and apply it to each query individually and independently. This flow is advantageous in its simplicity and generality. However, it may run into a performance hit in particular when successive commands exhibit similar properties where computation and data structures can be reused rather than starting from scratch. In our example, both query 1 and query 2 request the path to go through pin J. In query 2, we can potentially reuse the downstream data structure of pin J from query 1 to speed up the path generation process.

Iterative generations of timing critical paths with a specific cone of logic is a common workload in local optimization algorithms such as window-based timing-driven placement and region-constrained rip-up-and-reroute [2–4]. These algorithms iteratively transform the local landscape of the design, followed by timing queries with similar properties to evaluate their strategy or heuristic. To speed up the process, we proposed in this paper a general cache framework for efficient generation of timing critical paths. We summarize our contributions as follows:

- **A general cache framework.** We introduce a general and flexible cache framework for path-based timing analysis. Our cache scheme effectively considers the similarity among different queries. It can be easily turned on and off with little overhead to suit different optimization algorithms.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

DAC '19, June 2–6, 2019, Las Vegas, NV, USA

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-6725-7/19/06...\$15.00

<https://doi.org/10.1145/3316781.3317744>

- **Graph contraction.** We develop a graph contraction algorithm to reduce the size of the timing graph without affecting the path accuracy. As a result, our algorithm needs only a few essential points to find a timing critical path.
- **Efficient multi-threading.** Our cache framework enables efficient use of multi-threading. Each thread can have its own cache data structure independent of others and run the path search algorithm in parallel to improve the performance.

We evaluated our algorithms in the TAU 2018 Timing Analysis Contest benchmarks with golden reference generated by an industry standard timer [2]. The experimental results showed our algorithm can effectively improve the path generation runtime especially when local similarity exhibit among queries. We also demonstrated near-zero overhead in disabling the cache to make the timer run in normal mode. We achieved the best performance compared to the top performers in the TAU 2018 contest.

2 PROBLEM FORMULATION

We stick with the problem formulation of the TAU 2018 Timing Analysis Contest – *Efficient generation of timing reports from an STA graph with updated arrival and required times* [2]. Given a timing graph, the goal is to quickly and accurately answer the timing query that meets user-specified requirements.

2.1 Input Files

The input circuit design is described in the following industry standard formats:

- **Verilog (.v):** This file describes a gate-level netlist of the design.
- **SPEF (.spef):** This file describes the parasitic RC network for the design.
- **Liberty (.lib):** This file describes the cell characteristics of the design. There are two liberty files, early and late for minimum and maximum delay calculation, respectively.
- **Timing (.timing):** This file describes the arrival time of PI and clock pin, the required arrival time of PO and the clock cycle.
- **Command and Operations (.ops):** This file contains the set of `report_timing` queries and their options.

2.2 The Command: `report_timing`

The goal is to implement the command “`report_timing`” that reports a set of top worst-slack paths that satisfy the user-given requirements as follows:

- **-through pin:** The reported paths should go through this pin at either rising or falling edge.
- **-rise_through/-fall_through pin:** The reported paths should go through this pin at rising/falling edge.
- **-disable pin:** The reported paths should not go through this pin at neither rising nor falling edge.
- **-rise_disable/-fall_disable pin:** The reported paths should not go through this pin at rising/falling edge.
- **-nworst=1:** The number of top worst critical paths will be reported.
- **-cppr=<true/false>:** It specifies whether to consider common path pessimism removal (CPPR) analysis or not.

- **-mode=<hold|setup|both>:** It can be set to “setup”, “hold”, or “both” to specify which session of critical paths to report.

Figure 2: Overview of our program flow.

3 ALGORITHM

In this section, we discuss the details of our algorithm. The overview of our program flow is shown in Figure 2. We first create the timing graph and build up internal data structures from the input files. Then, we deal with each query in two steps. In the first step, we focus on identifying the search space from the specified through pins and our cache. Next, we demonstrate how to transform the problem into the graph search problem. We apply the algorithm in [3] to obtain the top- k worst post-CPPR slacks paths. Without loss of generality, we explain our algorithm only for the max (setup) mode. The same concept applies to min (hold) mode.

3.1 Preprocessing

3.1.1 Timing graph. We construct a directed acyclic graph (DAG), G , to represent the timing graph extracted from the input files. In G , each pin in the circuit has four nodes representing the combination of rising/falling transition and early/late mode. After building the timing graph G , we perform the forward and backward propagation to obtain the timing information, arrival time and required arrival time, of each node. The weight of each edge (u, v) in G is the delay from u to v . Additionally, we define $level(x)$ to be the topological order of node x in G . Obviously the four nodes of the same pin have identical topological order. We further define $level(pin)$ which equals to $level(x)$ where x represents the nodes of pin .

3.1.2 Preprocessing of `report_timing`. To incorporate each query into our cache, we discard the transition requirements of through pins; we modify their parameters, “-rise_through” or “-fall_through”, to “-through” and add its reversed transition type to the disabled group. For example, the query, “`report_timing -rise_through P1 -fall_through Pout`”, will be transformed to

“report_timing -through P1 -through Pout -fall_disable P1 -rise_disable Pout”. Moreover, we sort the through pins by their topological order in the timing graph, which lets us identify the search space easily.

3.2 Identify the Search Space for Each Query

We define G' as the search space which satisfies two conditions: (1) all possible valid paths are contained in G' (2) all paths in G' are valid, where a valid path means it must traverse all through pins but disabled pins. Additionally, let $G' = (V, E)$ be a DAG, where V and E are subset of vertices and edges in G . To facilitate obtaining the critical paths in the setup mode, the weight of $edge(u, v)$ in G' is the negative delay between u and v . We define $node(pin)$ as the nodes representing the *pin* and $through(i)$ as the i -th through pin. In G' , we use *startpoints* and *endpoints* to represent the nodes belonging to PI/FF:clk and PO/FF:d, respectively.

Intuitively, the search space G' consists of various independent path sets due to the property of a DAG. To be more specific, we define $paths(i, j)$ as the set of edges that connect paths between the two sets of nodes, i and j , where these edges exclude all disabled pins. Therefore, G' is equal to

$$\begin{aligned} S &= paths(startpoints, node(through(0))) \\ M &= \bigcup_{i=0}^{n-2} paths(node(through(i)), node(through(i+1))) \\ D &= paths(node(through(n-1)), endpoints) \\ G' &= S \cup M \cup D \end{aligned}$$

, where n is the number of through pins. An example G' of the query, “report_timing -through T1 -through T2”, is shown in Figure 3. For S and D , We simply perform backward and forward DFS from $through(0)$ and $through(n-1)$ to obtain them, respectively. For each term in M , $paths(node(through(i)), node(through(i+1)))$, we can also apply DFS from $node(through(i))$ to collect the edges to $node(through(i+1))$.

Figure 3: Search space G' can be divided by through pins.

3.3 Cache the Fan-in Cone of a Pin

To improve the overall performance and avoid duplicate calculation, we cache the fan-in cone of the most frequently used pins. For pin x , the upstream search space of x can be reused and applied to other queries also requesting to pass x . If we know the fan-in cone of $through(i+1)$ in advance, then we can largely reduce the time of obtaining $paths(node(through(i)), node(through(i+1)))$ by making sure that traversed nodes are in the fan-in cone of $through(i+1)$ when we perform DFS from $through(i)$. Therefore, finding the fan-in cone of all through pins is the necessary step to construct G' . We define a class, *CacheNode*, to store the fan-in cone for each pin. The essential members of *CacheNode* are listed below.

- *pin*: This *CacheNode* is attached to *pin*.
- *pre_pin*: The previous through pin of *pin* in current query, if *pin* is the first through pin then *pre_pin* is *None*.
- *is_in_fanin*: A table to record whether a node is in this fan-in cone.
- *starts*: It stores the set of nodes corresponding to startpoints in this fan-in cone.
- *searched_level*: The level of fan-in cone is searched. If the level is greater than $level(pre_pin)$ then extending the fan-in cone covered by the cache is necessary.
- *frontier*: The set of nodes to extend the fan-in cone.
- *last_used*: Last time of cache hit, the initial value is the time when it was created.
- *cache_hit*: The number of cache hits, the initial value is one.

A through pin is a cache hit pin if there exists a *CacheNode* for it or a cache miss pin otherwise. By the user settings, the maximum number of *CacheNode* would then be decided. Since the number of *CacheNode* is fixed, the cache replacement policy (Section 3.5) is required if there is no space to store the new *CacheNode* for cache miss pins.

LEMMA 3.1. For $through(i)$ in each query, its effective range of fan-in cone is between $level(through(i-1))$ and $level(through(i))$.

PROOF. For each query, the fan-in cone of $through(i)$ is only used for collecting $paths(node(through(i-1)), node(through(i)))$. During traversing from $node(through(i-1))$ to $node(through(i))$, we only pass the nodes which level are between $level(through(i-1))$ and $level(through(i))$ by the property of DAG. $\square$

Lemma 3.1 tells that it is not necessary to search all fan-in cone of a through pin when its *CacheNode* is created. When we initially search the fan-in cone of $through(i)$ and find that a node x is out of the effective range in the current query ($level(x) \leq level(pre_pin)$), we stop traversing and put x into *frontier*. If the fan-in cone of a *cacheNode* is too small for another query ($searched_level > level(pre_pin)$), we can simply extend fan-in cone from *frontier*. Figure 4 illustrates the fan-in cone of through pins in query 1 and query 2. In query 2, although pin J is a cache hit pin, its *pre_pin*, A, forces us to update the fan-in cone of J.

3.4 Condense the Timing Graph

Reducing the size of the timing graph can largely speed up the algorithm due to reduced search space. We contract the timing graph

Figure 4: The fan-in cone identified from query 1 and query 2. In query 2, we update the fan-in cone of J from its frontier, C and I.

after we construct it. Our principle is to recursively eliminate the node x which has only one in-degree and non-zero out-degree in the timing graph G . Under these circumstances, we define $par(x)$ as the fan-in of x , and conclude that x has condensed by the $par(x)$ by merging x into $par(x)$ and creating the edges from $par(x)$ to all fan-out of x .

LEMMA 3.2. *The input pins of the gates can be eliminated by our contraction method.*

PROOF. In the circuit, the input pins in the gates only come from the output pin of another gate; therefore, it has only one in-degree intrinsically. On the other hand, they also have non-zero out-degree as a result of connecting to the output ports in their gates. $\square$

Meanwhile, we record whether a node has been condensed by using the table *condensed_by* to indicate which node has condensed it. For element x in *condensed_by*, the initial value is x , which represents that x doesn't be condensed. If the query demands to pass the eliminated node x , then we replace x with *condensed_by*[x]; that is, we can know if the path needs to pass x , then it must pass *condensed_by*[x] first.

There are a few implementation details after condensing the graph. Although we have replaced condensed through pin x with *condensed_by*[x], we cannot keep correctness unless a check is placed on the edges starting from *condensed_by*[x] to avoid skipping the original through pin x . This stems from the fact that not all the edges starting from *condensed_by*[x] pass x .

Figure 5: An example of graph contraction. Query 1, "-through B -through J", is transformed to "-through A -through J" in the contracted graph.

An example of contraction is given in Figure 5. The gray nodes represent that they have been condensed. The final G is shown as blue lines. Query 1, "report_timing -through B -through J", will be modified to "report_timing -through A -through I" by *condensed_by*. Subsequently, we construct G' for this query. When traversing from node A , we observe that its fan-out, I and J , are both in the fan-in cone of the next through pin J . However, there exists a original through pin, B , between A and J , which entails us examining the edges starting A . For the edge (u, v) , we can tell if it is valid by checking whether there are some original through pins which level is between $level(u)$ and $level(v)$. If not, then this edge is valid in G' ; otherwise, we further check edge (u, v) whether the original through pin was condensed by it. Notably, after the condensing, there are two edges from A to I , edge (4) and (5). However, only the edge (4) is valid in G' since the other edge will make us leap the original through pin B . As a result, for those edges starting from A , only the edge (3) and (4) are valid in G' . Furthermore, this checking process is also executed for those disabled pins for avoiding picking the edges containing the disabled nodes.

3.5 Cache Replacement Policy and Overhead

We introduce a cache replacement policy to cope with the problem of finding the victim *CacheNode*. The settings in our cache scheme are listed as follows.

- *cache_size*: This value indicates the maximum number of *CacheNode* we store. In general, this value should be greater than the maximum number of through pins in all queries.
- *interval_time*: If a *CacheNode* has not been hit for *interval_time* queries, it has the high priority to be replaced.

Our replacement policy has two stages. The first stage is to select a *CacheNode* which has not been hit for *interval_time* queries. After that, if we still fail to find a victim, we select one from *cache_hit*. Intuitively, we gain the benefit from replacing the *CacheNode* having the minimum *cache_hit*. In our experiment, the *cache_size* and *interval_time* are set to 20, 10, respectively.

As for the memory cost for cache, it depends on the data structure to store the fan-in cone of a node. In our implementation, we use a bitset to identify whether a node is in the fan-in cone of a through node or not. Therefore, the memory overhead of each *CacheNode* is $O(n)$, where n is the total number of nodes in G .

3.6 Search the Critical Paths

Our final step is to obtain the k -worst paths from the G' containing all possible valid paths. If the query considers the CPPR, then we directly make use of the algorithm introduced in [3, 4] on G' to get the top k worst post-CPPR slack paths. Otherwise, we convert the problem of finding the top k worst slack paths into the typical k -shortest path problem. First, we define $path(s, d)$ as a path starting from a startpoint s to an endpoint d , and define $w(e)$ and $w(u, v)$ as the weight of edge e and edge (u, v) in G , respectively. In the late mode, the slack of a $path(s, d)$ referring to the slack of the endpoint d is equal to the $rat_d - at_d$, where rat_d and at_d represent the required arrival time and arrival time of d , respectively. Moreover, in the path-based analysis, the arrival time of an endpoint depends on the paths, the at_d of a $path(s, d)$ is

$$at_d = \sum_{e \in path(s, d)} w(e) + at_s$$

We thus add two nodes, *PseudoSource* and *PseudoDest*, to G' . After that, we connect *PseudoSource* and *PseudoDest* to *startpoints* and *endpoints* respectively and the weights of these edges are listed below:

$$\begin{aligned} \forall x \in startpoints, w(PseudoSource, x) &= -at_x \\ \forall x \in endpoints, w(x, PseudoDest) &= rat \end{aligned}$$

, where *startpoints* and *endpoints* are defined in Section 3.2. Consequently, the weight of a $path(PseudoSource, PseudoDest)$, sum of the weights of the edges that constitute it, is a slack since the weight of edges in the original G' is the negative delay between two nodes (Section 3.2). In the late mode, the worst slack path has the minimum slack; thus, the problem has turned into the s - t k shortest problem, and there are many state-of-the-art algorithms to solve k shortest path problem [5, 6].

3.7 Multi-threading

The structure of our algorithm allows us to easily extend it to multi-threading. For each thread, we maintain a private local cache. That is, each thread has its own sets of *CacheNode*. To search the path, we equally divide the endpoints to chunks and distribute the workload across threads. The final result is gathered through a parallel reduce operation.

4 EXPERIMENTAL RESULTS

In this section, we discuss the performance of our cache system and graph contraction techniques. We compare our timer with the winner of the TAU 2018 timing analysis contest. We will demonstrate our performance in large command set and scalability with increasing number of threads. We implemented our algorithm in C++ language on an Intel Xeon 2.4 GHz Linux machine with 32 GB memory which can run 16 threads concurrently.

4.1 Benchmarks

We use the TAU 2015 and 2018 contest benchmarks to evaluate our algorithm [2, 7]. The golden reference was generated by industry standard timers. In addition, we modified the contest benchmarks to incorporate more queries to demonstrate the scalability of our algorithm. A group of queries is generated with 4 to 5 common through pins (the startpoints and endpoints are included in the common through pins) and 2 to 3 different through pins randomly. Moreover, We define the similarity of a benchmark as the number of queries within a group. Following the rules of the contest [2], each query includes a startpoint and an endpoint. We query the most critical path for the setup check without CPPR.

4.2 Performance of our Algorithm

Table 1 shows the overall performance of our algorithm under different configurations. We compared the result with iTimer, the winner of the TAU 2018 contest. With graph contraction and cache, we outperformed iTimer across all benchmarks. In many cases, our graph contraction algorithm is able to reduce the graph size by more than $2\times$ (e.g., $2.4\times$ in *leon2_iccad*, $2.5\times$ in *tau2015_tip_master*). Enabling the cache allows us to reduce the runtime by about 2-20% (19% in *leon2_iccad*, 18% in *netcard_iccad*). The largest difference between our result and iTimer is observed in circuit *mgc_matrix_mult_iccad*, where we accomplished the queries in 45s but it took more than 10 minutes for iTimer to finish. Under this scenario, turning off the cache still results in faster runtime than iTimer (80%-85% faster). This purely demonstrated the effectiveness of our path-based analysis algorithm without the aid of cache.

Figure 6: Runtime comparison between our algorithm and iTimer under different similarity values (among 1M queries).

We next explore the effectiveness of our cache framework in increasing the similarity among subsequent queries. In this experiment, we focus on the two largest benchmarks, *leon2_iccad*, and *netcard_iccad*, and measure the runtime on different similarity values from one to ten. As shown in Figure 6, our timer is consistently faster than iTimer across all similarity values (123s versus 168s on average). We also observed similar trend when increasing the number of queries, as shown in Figure 7.

Table 1: Performance of the Proposed Algorithm on the TAU 2015 and 2018 Benchmarks and its Comparison to the TAU 2018 Contest Winner iTimer [2]

circuit	before contraction		after contraction		iTimer	with contraction			without contraction			TRD
	V	E	V	E		w/ CA	w/o CA	CRD	w/ CA	w/o CA	CRD	
leon2_iccad	8357749	6361534	3429588	2931946	187	123	151	19%	273	325	16%	62%
netcard_iccad	7802687	7068052	3902650	3165402	144	97	118	18%	237	266	11%	64%
leon3mp_iccad	6535987	5250406	2951790	2298616	99	75	87	14%	143	153	7%	51%
b19_iccad	1552641	1913614	1050296	863318	119	45	53	15%	76	81	6%	44%
vga_lcd_iccad	1324359	1166890	679900	486990	52	41	46	11%	74	79	6%	48%
mgc_matrix_mult_iccad	979341	1182124	696668	485456	>10 min	45	50	10%	67	73	8%	38%
mgc_edit_dist_iccad	889387	974964	630992	343972	251	52	61	15%	186	217	14%	76%
edit_dist_ispd	833219	974964	630992	343972	324	51	59	14%	178	220	19%	77%
vga_lcd	761461	673604	360908	312696	23	26	29	10%	36	39	8%	33%
des_per_ispd	743175	789476	476066	313410	30	30	33	9%	40	38	-5%	21%
cordic_ispd	255987	310320	182190	128130	146	47	51	8%	78	84	7%	44%
fft_ispd	232279	272408	153082	119326	25	29	29	0%	34	35	3%	17%
tau2015_tip_master	191049	140574	73598	66976	>10 min	85	90	6%	92	90	-2%	6%
tau2015_softusb_navre	39093	42440	21618	20822	28	29	30	3%	34	37	8%	22%
tau2015_cordic_core	19341	15932	9368	10164	80	45	46	2%	47	48	2%	6%

Note: Execution time is measured in seconds.

|V|: size of nodes; |E|: size of edges; w/ CA: with cache; w/o CA: without cache;

CRD: runtime reduction ratio from cache; TRD: total runtime reduction ratio;

Figure 7: Runtime versus different number of queries.

Figure 8: Runtime scalability with 1M queries under different number of threads.

4.3 Multi-threading Scalability

In this experiment, we demonstrate the performance of our algorithm under different number of threads. We consider three different similarity values on 3, 5, and 10, and measure the runtime using 1–10 threads. As shown in Figure 8, our algorithm scales with increasing the number of threads. With five threads, we are able to reduce the runtime by 3× in comparison to one thread.

5 CONCLUSION

In this paper, we have presented an efficient cache framework for the generation of timing critical paths. Our framework efficiently supports (1) a general cache scheme to capture the similarity among timing queries, (2) a graph contraction technique to reduce the search space without loss of accuracy, and (3) multi-threading. We have evaluated our algorithm on the TAU 2015 and 2018 contest benchmarks. On average, we can gain 10~15% speed-up from our cache scheme. Our graph contraction algorithm can reduce the search space by more than 2×.

6 ACKNOWLEDGMENT

This work is partially supported by NSF Grant CCF-1718883 and DARPA Grant FA-650-18-2-7843.

REFERENCES

- [1] J. Bhasker and Rakesh Chadha. *Static Timing Analysis for Nanometer Designs: A Practical Approach*. Springer Publishing Company, Incorporated, 1st edition, 2009.
- [2] ACM TAU 2018 Contest: Efficient Generation of Timing Reports From an STA Graph with Updated Arrival and Required Times, 2018.
- [3] T. Huang and M. D. F. Wong. OpenTimer: A high-performance timing analysis tool. In *2015 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*, pages 895–902, Nov 2015.
- [4] Y. Yang, Y. Chang, and I. H. Jiang. iTimerC: Common path pessimism removal using effective reduction methods. In *2014 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*, pages 600–605, Nov 2014.
- [5] Jin Y. Yen. Finding the k shortest loopless paths in a network. *Management Science*, 17(11):712–716, 1971.
- [6] David Eppstein. Finding the k shortest paths. In *35th Annual Symposium on Foundations of Computer Science, Santa Fe, New Mexico, USA, 20-22 November 1994*, pages 154–165. IEEE Computer Society, 1994.
- [7] ACM TAU 2015 Contest: Incremental Timing and Incremental CPPR, 2015.