

# Profiling over Adaptive Ranges

Shashidhar Mysore    Banit Agrawal    Timothy Sherwood  
Nisheeth Shrivastava    Subhash Suri

Department of Computer Science  
University of California, Santa Barbara  
{shashimc,banit,sherwood,nisheeth,suri}@cs.ucsb.edu

## Abstract

*Modern computer systems are called on to deal with billions of events every second, whether they are instructions executed, memory locations accessed, or packets forwarded. This presents a serious challenge to those who seek to quantify, analyze, or optimize such systems, because important trends and behaviors may easily be lost in a sea of data. We present Range Adaptive Profiling (RAP) as a new and general purpose profiling method capable of hierarchically classifying streams of data efficiently in hardware. Through the use of RAP, events in an input stream are dynamically classified into increasingly precise categories based on the frequency with which they occur. The more important a class, or range of events, the more precisely it is quantified.*

*Despite the dynamic nature of our technique, we build upon tight theoretic bounds covering both worst-case error as well as the required memory. In the limit, it is known that error and the memory bounds can be independent of the stream size, and grow only linearly with the level of precision desired. Significantly, we expose the critical constants in these algorithms and through careful engineering, algorithm re-design, and use of heuristics, we show how a high performance profile system can be implemented for Range Adaptive Profiling. RAP can be used on various profiles such as PCs, load values, and memory addresses, and has a broad range of uses, from hot-region profiling to quantifying cache miss value locality. We propose two methods of implementation, one in software and the other with specialized hardware, and we show that with just 8k bytes of memory range profiles can be gathered with an average accuracy of 98%.*

**Keywords:** Profiling Hardware, Range Adaptive, Value Locality.

## 1 Introduction

Many proposed run-time systems rely on profile information to make informed design and optimization decisions. Procedure and data placement, trace scheduling, value specialization, network load balancing, dynamic compilation, and a whole host of power management techniques can all be guided by an accurate picture of what a program is doing and how it is interacting with the system. A major problem in dealing with streams of profile data generated is that we can only store a small amount of information yet we need to

be able to accurately characterize the behavior of the entire stream. This is especially problematic if the profile information is not completely dominated by a small number of frequently seen, or “hot”, events.

A significant difficulty in gathering run-time profiles is keeping track of this data in a manner that requires little storage, incurs limited or negligible slowdown, and provides a consistent, accurate, and useful summary of the data. Profiling a large program for an extended amount of time (minutes or even hours), as required in a real system, results in the generation of huge amounts of data. Dealing with these large profiles in software requires clever schemes for adaptively sampling [21], compressing [43] and compacting [26] profiles to reduce the impact on memory.

The aim of our research is to explore a new profiling method capable of summarizing profile data in a streaming fashion (one-pass) with only a small and bounded amount of memory. *Range Adaptive Profiling* (RAP) uses a small set of counters to track *ranges* of profile data such as blocks of data and IP-addresses, segments of code, or ranges of load values. Every piece of data fed into the system is accounted for in some range (RAP merges the data rather than sample or filter), but the ranges which are chosen for profiling are adjusted dynamically based on observed program behavior. While perhaps not every type of profile can be merged easily into adaptive ranges, hot code regions can be found to guide optimization, ranges of values can guide encoding decisions and value prediction, while ranges of data memory will correspond to instances of data structures. Other types of profiles, such as edge profiling, can also be mapped onto adaptive ranges with simple extensions to the method. In Section 4 we describe three example uses of RAP in more detail.

In particular our paper makes the following contributions:

- We present the idea of Range Adaptive Profiling and show how it can be used to generate online summaries of different types of profile data including code, load values, memory content, and narrow-width operands.
- We describe Range Adaptive Profiling Trees and show how optimizing the branching factor and merging behavior can provide an implementable solution with guarantees on both summarization error and bounded memory.

- We present a method by which Range Adaptive Profiling can be efficiently pipelined if specialized hardware support is added.
- We quantify the Range Adaptive Profiling error and memory requirements for hot code regions and load value ranges. With as little as 8k bytes of memory, accuracy of up to 98% is possible.

The rest of the paper is laid out as follows: In Section 2 we begin by describing our online algorithm, while implementation details of our design are discussed in Section 3. In Section 4 we quantify the advantages of our scheme and provide some qualitative evidence of its usefulness in the form of range profiles. We describe related prior works in Section 5 and finally conclude in Section 6.

## 2 Profiling with Adaptive Precision

The first difficulty in building a run-time profiling system is in gathering the raw data. Several software techniques, such as binary instrumentation [5, 30, 37, 38] and sampling [2], can be used to generate and analyze this profile information with only a moderate amount of overhead [2, 3, 6, 7, 8, 14, 21, 26]. Recently, several researchers have proposed various forms of architectural support [1, 9, 10, 13, 18, 31, 32, 34, 40, 47] with the aim of increasing accuracy and further reducing the overhead of software based techniques. Value profiles can be exploited to perform code specialization [6], value prediction [28, 45], and value encoding [40, 41]. Operand profiles identify the potential to apply power and performance optimizations [29, 4]. Address profiles have been used for data layout optimizations [33] and data prefetching mechanisms [8], and code profiling for focusing optimization efforts on the most important regions of a program. Control flow traces and path profiles [3, 26, 44] can be used to perform path sensitive optimizations [17, 42] and path sensitive predictions [23]. A general purpose framework for dealing with profile data has even been proposed [43]. While gathering data is a difficult problem, it is not the end of the story.

To explain the concept behind range adaptive profiling let us start with a simple example. Suppose we would like to know something about the regions of code that *gcc* is spending its time in. The simplest and lowest precision way to quantify this is to keep one counter which counts all instructions executed on behalf of *gcc*. The counter keeps a perfectly accurate count, and covers the entire program, but of course the profile has *no precision* and fails to provide any information on which subset of instructions is really the most important. If two counters are available, the next logical step might be to have one count the “top” half of *gcc* code and to have the second counter track the “bottom” half of *gcc* code. In this example each counter is tracking a range of code in *gcc*, although as we discuss in Section 4, this works equally well for memory addresses, values, and other range based profile types. This idea of dividing the code into  $N$  ranges for  $N$  counters could be easily extended to 4, 8, 16 counters and so on. Unfortunately, this quickly gets out of hand, and to track

the program at the precision of an instruction we would need counters for each and every basic block.

Our technique is based upon the realization that not all profile information is equally valuable. The more frequently a set of events occurs, the more important it is to precisely quantify and characterize this set of events. Specifically, it may be sufficient to group profile data into *ranges* - where the most frequently occurring ranges of events are identified and broken into more precise ranges while the least frequently occurring events are kept as larger ranges. If the profiling ranges are properly managed over time, we can strike a balance between profile resolution and overhead.

When a particular range of events constitutes a significant portion of the total profile, then that range should be subdivided and profiled more precisely. This recursive refinement of profile ranges maps nicely onto a tree, where the root of the tree represents the entire range of events and each child of a node represents a refinement of the profiling for a particular subrange. We formalize this idea as *Range Adaptive Profiling* and show how we implement this idea in a specialized hardware scheme.

### 2.1 Profile Trees

To gather profiles where the granularity is changing dynamically, we will need a data structure in which we can store our profiles. The majority of the past work in this area has assumed a flat storage of the profile. Whether the data was gathered through hardware performance counters [1], stratified sampling [34], or even potentially in fixed ranges [47, 46], the end result is essentially a list of equivalent items and their counts. While there exists some specialized software and hardware systems that attempt to tightly compress particular types of traces [1, 9, 10, 31, 34, 47], we believe we are the first to present a general hardware-based methodology for storing profiles in a *hierarchical* fashion.

As we mentioned above, the most natural way to store our hierarchical profiles is with a tree. This tree will keep a constantly up-to-date summary of the data stream, and in this section we describe the three types of operations on the tree that we need to support. The first, and by far the most common, operation is a simple *update*, where a counter in the tree is simply incremented to track the incoming data. To refine the granularity of a sufficiently hot range and to ensure precision, we have a *split* operation. Finally, we need *merges* to combine together relatively unimportant data which ensures that the tree is carefully pruned to maintain the least number of counters necessary to capture all the important information. The splits and the merges change the structure of the tree and hence dynamically re-map profile events to the counters. While at a high level this simply sounds like a simple tree, in reality each of these three functions has been specially designed such that the overall data structure is both implementable in an online and pipelined way and provides a type of worst case bound on error (the  $\epsilon$ -error discussed in Section 2.2). While all three operations are discussed in detail, we begin with a discussion of the simple update.

The profile tree is built of nodes, and each node corresponds to a particular *range* of events that the profiler might see. As was mentioned earlier, the root node represents the entire range of events possible, and each child of a node will capture a proper sub-range of its parent. When an event enters the profiling system, for example the PC of a cache miss, this event is matched into the range that covers it. Because an event often matches several possible ranges, we need to find the *smallest* range that includes that event and then increment that node’s counter. This will ensure that profiling is done with as much precision as possible without modifying the tree. As an example, please refer to Figure 1. This figure is a snapshot of the tree structure with each node in the tree tracking a range of values  $[min\_range, max\_range]$  and a *count* to track the number of times an event entering the profiler mapped to this node as the smallest range covering the event. (Split and Merge operations referred to on this figure are explained in the following subsection). If an incoming event had a value of 12, in the graph on the left hand side, it would match the ranges of  $[0, 255]$ ,  $[0, 63]$ , and  $[12, 13]$  but only the node responsible for profiling  $[12, 13]$  (smallest range) would have its counter updated.

## 2.2 Growing the Trees

While updates are the most common operation performed on a tree, updates do not modify the structure of the tree to adapt to the input stream. The two operations that actually modify the tree structure are *split*, which further refines the profiling of a given range, and *merge*, which decreases the granularity with which a range is profiled. The main idea behind a split is that if a range is important enough, its counter will increase faster than average. Eventually this node will grow so ripe, that it makes sense to burst the node into a number of subranges. In this way the tree grows to increase precision where the profile has more weight. If the counters from a set of ranges are no longer a sufficient portion of the whole stream they can be merged together with little impact.<sup>1</sup>

The key to applying the split and merge operations is knowing *when* they should be applied. This can be done by setting a *SplitThreshold*, and any node that grows larger than this threshold should sprout children to more accurately profile each of the sub-ranges. The *SplitThreshold* is a function of the number of events processed  $n$ , and the maximum possible height of the tree  $\log(R)$  where  $R$  is the maximum range to be considered. Specifically we set

$$SplitThreshold = \frac{\epsilon \cdot n}{\log(R)}$$

where  $\epsilon$  is a user defined constant between 0 and 1. If the *SplitThreshold* is set in this way, the *maximum* amount of error possible, relative to the entire input stream<sup>2</sup>, is  $\epsilon$ . For ex-

<sup>1</sup>Counters are never decremented which is why this is not a sampling scheme, rather merges happen when the rest of the tree has outgrown a particular set of regions.

<sup>2</sup>The  $\epsilon$  error is defined as a fraction of the total length of the input stream, while percent error is error relative to the actual count of a range

ample, if the user sets  $\epsilon$  to 1%, that means for any given range the estimate for that range will never be off by more than 1% of the total events processed. Further, it can be shown that the maximum amount of memory required by a tree built with this split threshold is  $O(\log(R)/\epsilon)$  [19]. The exact byte counts, overheads, and percent errors are described in Sections 3 and 4.

**Split** - Calculating when an update needs to be followed with a split operation is actually a fairly straightforward task. We simply compare the value of the counter with the split threshold described above. Any time a node grows over this limit, we need to add a set of children to this node that cover and subdivide its range. The original node keeps its counter, and each of the children have their counts initialized to zero.

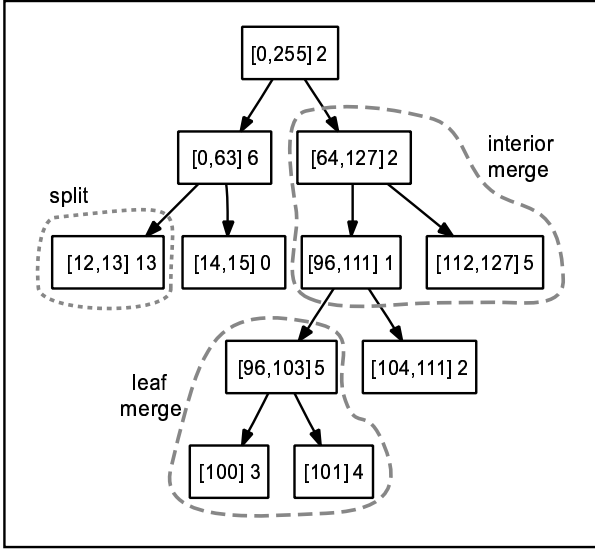
**Merge** - While splitting is crucial to the adaptation of the granularity, if all we ever did was update or split, it would be impossible to bound the total amount of memory required. For example, a region of code may start out “hot” and as such might have been split into many separate counters to count every basic block within. Later, however, it may turn out to be relatively unimportant and we may wish to release all counters associated with the basic blocks and retain just one counter for the entire region. A way to un-split a set of regions is to do a merge operation. Rather than simply throwing away the range profile information from each of the children nodes, we incorporate them into the parent node. Because any count gathered for a child is equally valid to be stored on the parent range (because it is a super-range), we simply sum together the count of the child nodes and add them into the parent.

## 3 Implementation Details

In order to build an effective profiling system around the algorithm described in Section 2, there are several tasks that need to be performed at runtime. First off, we need a mechanism to gather profile data. In a purely software based approach, these profiles can be generated through either binary instrumentation [5, 30, 38, 37] or hardware performance counters [1, 11, 12, 20, 22]. Even in a software based approach the input data should be buffered to some extent and duplicate values should be merged together to help improve performance. In the case of a hardware-assisted or hardware-only approach, we assume that the profiles are generated using a pre-existing or proposed profiling structure [9, 13, 18, 47]. Specifically, we assume a structure similar to ProfileMe [13] for collecting the input events (load values, PC, memory addresses, etc.). The buffered events are processed one after another in the order the load instructions (for value profiling) or branch instructions (for code profile) retire. The buffer size and the sampling module will affect the overall accuracy of a profile, but it has no impact on the way in which Range Adaptive Profiling summarizes the data and for this paper we concentrate solely on the accuracy of the summarization step.

Processing the gathered events and maintaining the RAP tree based counter structure can be time consuming if implemented naively. In this section we discuss ways of speeding

RAP Tree before Merge Cycle (split threshold=13)



RAP Tree after Merge and then Split

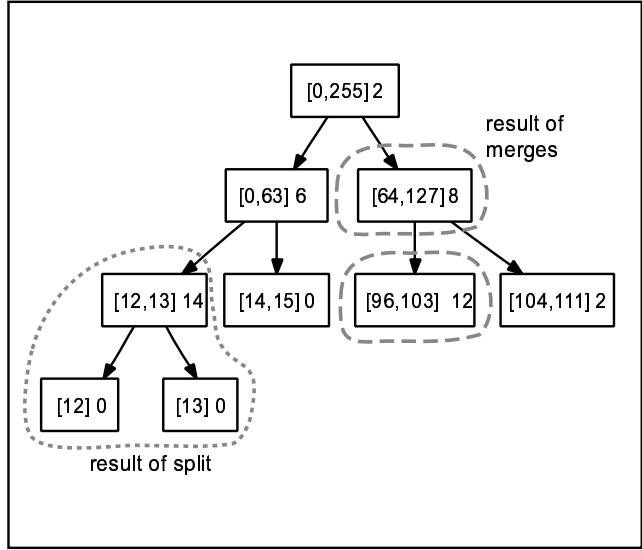


Figure 1: A range adaptive profiling tree (in this example, each node has 2 out edges). The diagram on the left is the state of the tree just before a merge cycle begins. During a merge cycle, the tree is walked and any set of nodes that have insufficient weight to warrant separate profiles are merged (in this example the cutoff is a cumulative weight of 13). Following this merge, an access to item 12 might occur, which will push the node that captures the range [12,13] to go over the split threshold. This would cause the node to be split into two different nodes, and subsequent accesses to item 12 or 13 will be recorded on an item by item basis.

up these tasks and describe the software implementation of RAP and also a hardware based approach that operates with little or no support from software.

### 3.1 Algorithm and Architecture Design Issues

To enable efficient storage and searches on the profile tree, a suitable *branching factor* ( $b$ ) must be used. The branching factor is the number of children that will be generated in a split operation. If  $b$  is too small, the ranges marked for profiling will take longer to converge on the best set. For example, if one particular value in a range is accounting for 100% of the profile data seen, it will take exactly  $\log_b(R)$  splits to finally start profiling this item individually which in turn effects the error in the profile. On the other hand, if  $b$  is too large, the amount of memory required to store the tree will grow. The higher the branching factor, the more extraneous children will be kept around. To seek a balance between these two constraints, we analyzed the effect of branching factor on the worst case number of nodes that can appear in the tree. Figure 2 shows this tradeoff. On the  $x$ -axis we have a variety of different branching factors, and on the  $y$ -axis is the worst case number of nodes that could be generated for a branching factor of  $b$  and an  $\epsilon$  of 1%. We found that a branching factor of 4 provides a good tradeoff between the required amount of memory and the effect on performance and error. Note that the figure shows *worst case* number of nodes and as will be shown in Section 4, in the common case the number of nodes is a factor of 1000 less.

Another problem that shows up in an implementation is when to perform the merges. Finding a node that needs to be split is easy, a counter is updated and then we check if

the counter is over the threshold. Finding the places where a merge must be performed is much more difficult, as they, by definition, happen away from where the updates are occurring. How does one detect when a merge is needed? One approach is to build a secondary merge heap, which stores a list of those nodes that are most in need of merging. While this approach is suitable from a theoretical standpoint, updating the merge heap requires many extra tree operations and a full additional tree. Furthermore, one merge can result in a new node which in turn needs to be merged into its parent and so on. Rather than detecting and handling merges at the soonest possible time, we propose batching the merges together.

By performing merges periodically, instead of in a continuous manner, we avoid the problem of having to continuously search the tree for valid sets of nodes to be merged. In order to grow, the tree must split, and in order to split, the counts of the nodes must grow past the split threshold. The key point to see is that as the number of events processed grows, the relative rate at which the tree can split *must* slow down. In fact, an un-merged tree can grow at a rate which is at most *logarithmic* with the number of events processed. Instead of having a fixed period for updates, we can have updates with an *exponentially increasing period* and the worst case bounds will still hold. This idea can be seen most clearly in Figure 3.

In Figure 3, the  $x$ -axis shows the number of events processed (instructions executed, values profiled, etc.). The  $y$ -axis is the worst case bound on the number of nodes required to profile with an  $\epsilon$  of 1%. At the beginning, and after every merge, the worst case number of nodes is bounded to  $384k$ . After a merge, the worst case size of the tree grows slowly, inching up at a logarithmic rate. If we wait for some num-




Figure 2: The two independent graphs plotted show the memory size requirement for different branching factors  $b$  (lower graph) and merge-interval ratios  $q$  (upper graph). We choose  $b = 4$  as it is a better tradeoff between memory consumed and the height of the tree. With  $q = 2$  we see that the memory size is the least.

ber of events  $e$  to pass before doing a merge operation, the next time around we can wait a total of  $2e$  events before the worst case number of nodes grows to the same point. In other words, if it took  $e$  events to force a split in the first period, in the second period the tree will be twice as big and it will require twice as many events in order for a split to be necessary. While in this example, we double the interval between consecutive merges, in general, we could increase the interval by a factor of  $q$ . In figure 2 we show that doubling the intervals is the most cost effecting setting for  $q$ .

### 3.2 Software Range Profiling

Using the results from the previous sections, we have developed a software implementation of RAP that can be called from software-only systems. The C++ implementation has an API with three methods `rap_init()`, `rap_add_points()`, and `rap_finalize()` which can either be called from online analysis or to post process trace files. `rap_init` initializes the RAP tree with an initial set of counters and appropriate range values. `rap_init` also initializes data structures to enable profiling multiple events simultaneously. The RAP tree is a dynamically allocated tree and `rap_add_points` looks up the appropriate counter, updates the counter, and when needed calls the internal functions `rap_split()` and `rap_merge()`. The post processing phase of deriving statistical inferences about the stream can be done through `rap_finalize`, it also dumps the resulting RAP tree in ascii format for further processing such as identifying hot-spots, range coverage, phase identification, and so on. Our software version is available at: <http://www.cs.ucsb.edu/~arch/rap>

### 3.3 Hardware Support for Range Profiling

While a software based approach has many applications, we are interested in using this technique at high speed in run-time systems. For example, we could use this method to analyze a front side bus trace to see what memory is being accessed when a program runs for minutes or hours. We have purposefully designed the algorithms to allow for an efficient and high speed hardware implementation which can tap

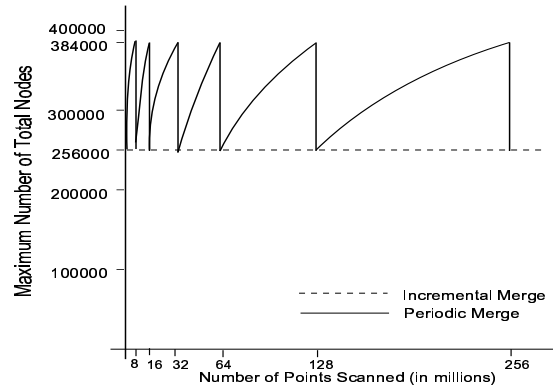



Figure 3: If merges are performed continuously, the tightest bound on the number of nodes required can be maintained but at the cost of continuously searching for merge opportunities every single cycle. Instead, we can still have bounded worst case memory requirements if we batch the merges together with an exponentially decreasing frequency.

into any streaming source of profile information. This source can be from a bus or debug port [16] off-chip, or plugged into the back end of any number of proposed on-chip profiling schemes. We are currently developing an FPGA prototype that can be interfaced through a high speed network or PCI-X. In this Section we describe a hardware design, and in Section 3.4 we quantify the design in terms of performance, power, and area.

The main features of our design can be seen in Figure 4. The profiling engine is divided into 5 main stages. In stages 0 and 1, the input events are first buffered, and then all matching ranges are found. In stage 2 the smallest matching range is determined, which then causes the appropriate counters to be updated. Splits and merges are special cases and require pipeline stalls. With the exception of the actual counter increment, each update to the profile tree is independent from the previous. Splits and merges require more work, because they create inter-event dependencies that must be satisfied before more events may be processed. However, compared to updates, splits and merges are very small in number, hence they have little impact on the performance and the total number of stalls is small and bounded.

**Stage 0:** The small buffer shown at stage 0 in Figure 4 stores incoming points. When the need to perform a merge occurs (periodically and at exponentially decreasing frequency), the pipeline is stalled while the counters are searched for potential merge sets. During this time events will stack up for ten to a hundred cycles, and we will need to keep them in a temporary storage so they can be processed later. In the case of a split, the pipeline will need to be flushed and reset to the point directly before where the split should have occurred. In this case the buffer will re-enter those events into the pipeline. It is quite possible to make this buffer pre-process the points by combining *identical* events. We have observed that a 1k buffer can reduce the throughput requirements on RAP by a factor of 10 for code profiling.

**Stage 1:** For every point fetched from the buffer, we need to find the set of ranges that include that point. This operation is very similar to the Longest Prefix Match and can be carried



**Figure 4: Architecture for the Pipelined RAP Engine.** *Stage0* - Shows the Event Buffer which buffers incoming events and the number of times it is seen since it was previously processed by the RAP engine. *Stage1* - A TCAM range matching provides indices to all TCAM cells which match the incoming event identifier (IP address, memory, PC, value, etc). *Stage2* - A priority Arbiter chooses the longest match by giving it the highest priority. *Stage3* - A set of counters are maintained, each corresponding to an entry in the TCAM array; priority arbiter chooses the counter to be incremented for the incoming event. *Stage4* - For every counter incremented in *Stage3*, a comparator checks the counter value against the current value of the split threshold and if necessary initiates a node split to adapt the precision of the profile maintained.

out in constant time with a Ternary CAM as shown in Figure 4. The TCAM sets the appropriate match line high, for all ranges that match. In order to figure out the *smallest range* which is also the *longest prefix*, the TCAM entries have to be partially sorted by prefix length. There can never be matches from two different entries of the same range width. Hence, this stage can be further pipelined by looking at nibble or byte for each comparison [27].

**Stage 2:** After the potential matches are identified, we need to find the longest prefix match, which should correspond to the last matching entry. Given  $N$  match lines in order, sorted by prefix length, finding the longest match is simply a matter of giving highest priority to longest matches and allowing only one match to proceed. This is exactly the function of a fixed priority  $N \times 1$  arbiter. The output of the highest priority line will trigger the word line of the matching counter. Note that while in this paper we assume a TCAM based approach, with a branching factor of  $b$ , the tree is really a *multibit trie* and there are a variety of techniques that can be used to build high speed implementations from network algorithms [36].

**Stage 3:** Once the smallest range match has been found, we simply need to update the appropriate counter. To handle a continuous stream of data to the array, one read port and one write port is needed.

**Stage 4:** The final stage compares the result of the updated counter with the split threshold. If the counter is above the split threshold then the node is expanded to have four children (for *branching\_factor*( $b$ ) = 4), each initialized to a zero count. The split and merge thresholds are stored in separate registers and recomputed whenever the number of events ( $n$ ) change. This computation can be done in parallel with other operations as it depends only on  $n$  and some predefined values. In our implementation, the split and merge thresholds

can be the same, hence just one computation and one register is sufficient. If a split is encountered, the pipeline may need to be flushed to properly account for these new nodes.

In our implementation a split requires making new entries in the TCAM and SRAM data array. Four new children nodes are created and inserted in the TCAM with the ranges set appropriately, covering a quarter of the parent range. Corresponding entries in the memory are inserted storing the counter and other information of the newly created nodes. A split node could be either a leaf node or a parent. If the node is a leaf then the split operation involves just setting of a pointer from the parent to the newly created children. If the node is already a parent, but its children do not cover the entire range of the parent (This could be the case after an internal merge as described in Section 2.2), then the split also involves an extra operation of identifying the new parent of the existing children and setting the children pointers. In terms of performance, these splits are not a large problem as there can be at most 6400 of them in a given interval, in our implementation.

A merge operation is even more expensive compared to other operations, but by batching them together we reduce the overhead significantly. Batch merges are initiated periodically and in every batch of merges entries in the TCAM are scanned bottom-up to find candidate nodes to be merged. Corresponding SRAM data array entries are then deleted. This recursive operation prunes the RAP tree to provide compacted profile information. If there are 4 billion ( $2^{32}$ ) events to be profiled, and we assume that there will be at least a thousand ( $2^{10}$ ) events before we do our first merge, then there will only need to be  $32 - 10 = 22$  different batches of merges. Similarly, to profile  $2^{64}$  events, requires  $64 - 10 = 54$  batches of merges. If we wish to profile large amounts of data, any cost of doing a merge is quickly amortized.

### 3.4 Analysis of Required Hardware

In this subsection, we estimate the power consumption, area and delay of various hardware components. We extracted and modified the power models from Cacti-3.2 [35] and Orion [39] tools to model the major components in our hardware design such as Ternary Content Addressable Memories (TCAMs), SRAM data array, comparator, priority arbiter, and registers. We then validated our results against published results from high speed circuit design conferences. We assume a very conservative  $0.18\mu\text{m}$  technology and we change the voltage supply and various other device parameters accordingly. In particular, we present the worst-case maximum delay and energy consumption by modeling the maximal switching for any particular operation.

Using 4096 x 36 TCAM and 16KB SRAM data array configurations and summing up the area of all hardware components we find that our Pipelined RAP Engine requires  $24.73\text{ mm}^2$  of area. The clock frequency is determined by the maximum delay in any pipeline stage and we find that it is governed by TCAM look up stage. The critical path delay in TCAM lookup stage is  $7\text{ ns}$ . We can aggressively pipeline the TCAM stage by doing byte/nibble comparison at each pipeline stage [27] and effectively we can shift the critical path to the SRAM stage, which takes  $1.26\text{ ns}$  time. We also add up the maximum energy components of all the hardware components and we find that a total of  $1.272\text{ nJ}$  energy is consumed. It is also true that an implementation of RAP that can handle 4k different ranges is very aggressive and would most likely be applicable for off-chip profiling, but that for a 400-node version the area and power would be more than a factor of 10 times less. On an average, RAP requires 4 cycles to process an event, and requires 2 cycles each for TCAM and SRAM accesses per event.

## 4 Evaluating Range Adaptive Profiles

In the sections leading up to this, we have presented the algorithms and designs necessary to perform range adaptive profiling. In this section we analyze the results of our effort by quantifying the memory requirements and errors involved across several SPEC benchmarks, and describe several example use scenarios.

### 4.1 Profiling with RAP

In trying to examine ranges in the code, values, addresses, or other parameters of a running program, RAP should focus in on the *hot ranges*. A range is considered hot if and only if the total count for that range and all its *non-hot* sub-ranges is above a certain threshold. Note that our definition excludes the possibility that a range is considered hot simply because it has one or more hot children. This is very useful because with a small fixed number of hot ranges we can accurately paint a picture of the distribution of events across the entire range of possible events. For example, when running RAP on a trace of basic blocks, our technique will automatically focus in on the most important regions of code, yet it will provide a balanced overview of the code as a whole. For `gcc` we iden-

tify seven distinct regions of the program where each region accounted for more than 10% of the instructions executed.

In addition to code profiles, we also wanted to truly demonstrate the abilities of our scheme by profiling a set of events that has significantly less locality than code profiles. While it has been shown that a single value may account for the top 20% to 40% of all load values, there is a large tail to this distribution which will stress our range profiling system. By building a RAP tree over the set of all values loaded by a program, we can calculate the ranges of values which would cover 50%, 80%, or even 95% of all loads. Figure 5 shows exactly this information for `gzip` and identifies all ranges of load values which are more than 10% hot. In this figure, the hot ranges of load value are shown (with min and max), and they are annotated with their relative weight.

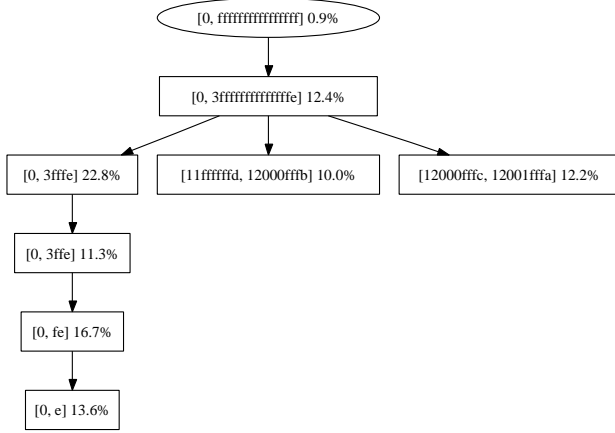
From this figure one can easily see that for `gzip`, load values in the range of `[0,e]` account for 13.6% of all loads, while the range `[0,fe]` *excluding* `[0,e]` accounts for 16.7%. Thus the entire range `[0,fe]` (including the hot sub-range) accounts for  $13.6\% + 16.7\% = 30.3\%$  of loads executed. These summaries are computed completely online and in hardware and could be used to guide optimizations such as value range specialization or to assist in value prediction.

For any profiling system to be feasible, the theoretical and empirical error and memory overheads need to be low. A theoretical analysis of RAP's memory use and error was overviewed in Section 2, and in Sections 4.2 and 4.3 we reevaluate these in the context of code and value profiling. We run our system on a set of programs from the SPEC benchmarks to completion, for reference inputs. The choice of these two types of profiles was governed by factors which can stress-test the RAP system. The locality present in code profiles will stress the upper bounds on memory required for RAP. The heavy tailed distribution of value profiles exercises the range adaptation aspects of RAP. In the rest of this section, we present an analysis of RAP with respect to memory required and error, and illustrate advanced profiling applications of RAP.


### 4.2 Memory Analysis

As explained previously, range adaptive profiling stores profiles hierarchically in a tree structure (RAP tree). The number of nodes in the RAP tree will tell us the memory requirement of this scheme. This section gives some practical estimates for various benchmark programs with each node requiring about 128 bits of memory.

Figure 7 shows different benchmarks on the  $x$ -axis and the maximum and average number of nodes required by RAP in evaluating these benchmarks is shown on the  $y$ -axis. The left hand two graphs show the maximum and average memory required for various benchmarks in identifying hot regions of a code for  $\epsilon = 10\%$  (top) and  $\epsilon = 1\%$  (bottom). As the tree grows between merge intervals and shrinks after a merge, the maximum memory is the largest of the tree sizes just before the merge operations during the entire run of a benchmark and the average number of nodes indicates the common-case



**Figure 5: Hot ranges among the load values in `gzip` as identified by RAP with  $\epsilon = 1\%$ . We see that there are 7 hot ranges which were encountered for more than 10% of the entire load value stream. Note that this tree is a subset of the RAP tree, showing only the hot nodes**



**Figure 6: Number of nodes required to track the basic blocks of `gcc` with  $\epsilon = 10\%$ . While the number of nodes is far less than the worst case bounds that we estimated, the pattern of growth (due to splits) and rapid reduction (due to batched merging at points marked by dashed lines) can be clearly seen.**

memory requirement. The two graphs on the right present similar parameters for value profiles. We see that a maximum of 500 nodes is sufficient to evaluate code profiles with  $\epsilon = 10\%$  for the set of benchmarks. In Section 4.3 we show that with this many nodes we can guarantee 98% accurate information. We can also observe that `gcc`, which has the highest number of distinct basic blocks, requires a maximum of 453 nodes in the RAP tree for code profiling. The graphs on the right of Figure 7, show similar trends for value profiling. `parser` which has the largest number of load values requires a maximum of 733 nodes and an average of 203 nodes in the RAP tree (Figure 7 for  $\epsilon=10\%$ ). Similarly, the RAP tree requires an average 300 nodes to provide 99% accurate information on load profiles.

An important observation to make is that RAP judiciously allocates counters only if it is sure it is worth allocating them. For example, since the locality among values is less, value profiling with RAP uses less memory (average 300 nodes) compared to code profiling (average 450 nodes) which has more locality. This advantage of being able to provide such accurate information using a small amount of memory, is attributed to the splits and merges we do on the RAP tree (as described in Section 2).

Back in Figure 3 we described the bounds on memory requirements as they change over time. To test what happens

in a real implementation, we generated Figure 6 which shows the variations of tree size for one such run of `gcc`. The  $x$ -axis represents the number of basic block vectors seen and the  $y$ -axis is the number of nodes in the RAP tree. We see a similar pattern to the theoretical expectation, which is the slow building of memory marked by periodic merges which maintain the overall bounds on resource consumption.

### 4.3 Error Evaluation

While the theoretical bounds on error are very useful, if our device is to be used to characterize dynamic program behavior in a real operating environment, the average and worst case percent errors<sup>3</sup> are extremely important.

Due to the way the algorithm is designed, the counts for a range in the tree is always a lower bound on the actual count. Hence, if RAP identifies a node as hot, then that node is guaranteed to be hot. A hot node means that a set of individual events in that range is hot. In cases where the range is a single event, we have identified a hot range with most precision<sup>4</sup>.

The split threshold is set in such a way that as soon as a node counts events more than a proportion of the total events seen, the node splits into sub-ranges. A merge, similarly, never merges ranges which are hot enough to warrant precise profiling. This ensures that RAP *always profiles with the smallest ranges possible*. Hence, for a given  $\epsilon$ , we can guarantee that RAP always identifies all hot ranges with the greatest precision possible.

Not only is it important to identify the most frequently observed ranges in a profile stream, but it is equally important to measure how accurately these ranges are quantified. For every hot region identified by RAP, the estimated counts of the events that contributed to the hot regions were used to compute the percent error. The numbers presented in Figure 8 are a comparison of the estimates that RAP made online, with the actual count that was gathered by making multiple passes through the program’s execution, tracking one hot range at a time (as a perfect offline profiler would). Figure 8 shows the percent error in estimating the counts on the hot ranges for each of the different benchmarks. *Maximum\_10* and *Maximum\_1* is the maximum of the percent errors among all the hot regions for a benchmark, identified in a RAP tree with  $\epsilon = 10\%$  and  $\epsilon = 1\%$  respectively. Similarly *Average\_10* (*Average\_1*) is the average of the percent errors for all the identified hot ranges within a benchmark with  $\epsilon = 10\%$  ( $\epsilon = 1\%$ ). The  $y$ -axis in Figure 8 shows percent error for various benchmarks. The graph on the left is a measure of accuracy when identifying hot regions of the code and the graph on the right shows different errors for load value analysis.

In the graph on the left in Figure 8, the benchmark `gcc` shows the highest maximum percent error of 13.5% with  $\epsilon =$

<sup>3</sup>Percent error is error relative to the actual count of an event, whereas  $\epsilon$  is the error with respect to size of the entire stream.

<sup>4</sup>By precision we mean the ability to zoom into profile ranges as narrow as possible, and by accuracy we refer to error in the quantitative profile information estimated by RAP with respect to a perfect profiler. A perfect profiler is one which can gather event counts with 100% accuracy




Figure 7: The number of nodes in the RAP tree which is an indication of the memory required by our profiler is plotted on the  $y$ -axis. The left hand two graphs show the maximum and average memory required for various benchmarks in identifying hot regions of a code for  $\epsilon = 10\%$  (top) and  $\epsilon = 1\%$  (bottom). And the two graphs on the right present similar parameters for value profiles.

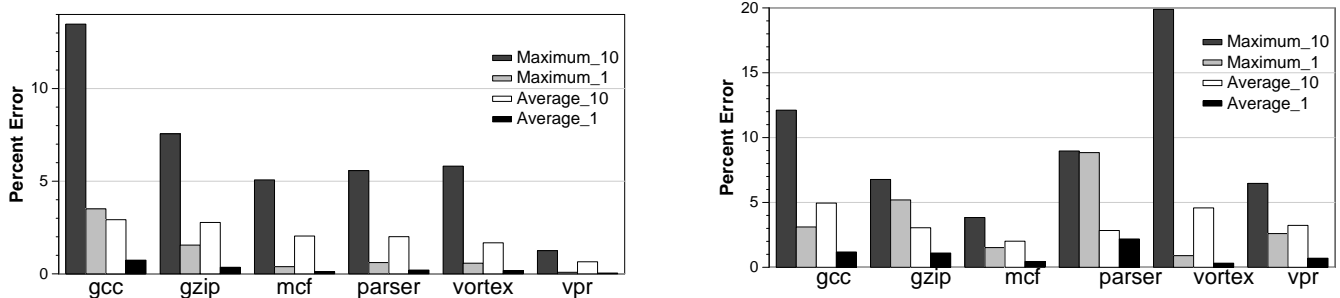


Figure 8: Percent Error for the hot events identified by RAP for various benchmarks is shown in this figure (for  $\epsilon = 1\%$  and  $10\%$ ). The graph on the left is a measure of accuracy when identifying hot regions of the code. The graph on the right shows similar values for load value analysis.

10%. This error of 13.5% was from a hot-range of the code, which was quite narrow and deep in the RAP tree, however, excluding this hot-range, the second maximum percent error in `gcc` is just 3.1%. An important point to draw from this graph is that with  $\epsilon = 10\%$ , the average percent error is still just about 2%.

Load value analysis, however, was more complex than code profiling because of the wide range of values within which incoming load events could be. With load value analysis (graph on the right in Figure 8), we see that `vortex` has the maximum percent error of around 20% which was due to the hot-value 0 (note, however, that this is still less than 10% error with respect to the entire stream). We also see a negligible percent error with  $\epsilon = 1\%$ ; and with  $\epsilon = 10\%$  an average of just 3.4% over all benchmarks. As can be observed, on an average RAP can provide 98% accurate information about code profiles and is 96.6% accurate on value profiles. Trends about program behavior, hot regions, value distribution, memory access patterns are some of the characteristics which can be easily and accurately detected with RAP.

To build a useful and feasible profiler, the error and memory requirements should be bounded absolutely, without reference to the stream length and the type of profile being analyzed. As we have just seen in this subsection, RAP not only precisely identifies range information on a stream of profile events efficiently, but also provides very accurate information.

#### 4.4 Additional Applications of RAP

Thus far we have discussed how RAP can be used to track code and value profiles, and use these to stress-test and evaluate our system. Here we describe several different scenarios where the capabilities of RAP would be useful including: cache-miss value profiling, narrow-width operand profiling, zero-load memory ranges.

**Cache-Miss Value Profiling** – While we have shown how RAP can be used to profile value locality in a more general sense than simply quantifying “hot values”, architects typically need to target cache misses, rather than simply all loads. Some have hypothesized that while value locality might be present, it may be greatly diminished when only the cache misses are examined. By simply building a RAP tree over the




Figure 9: This figure shows how RAP can be used to extract insightful information about value localities. The  $x$ -axis shows the number of bits required to represent the hot ranges of values and  $y$ -axis represents the percentage of values profiled

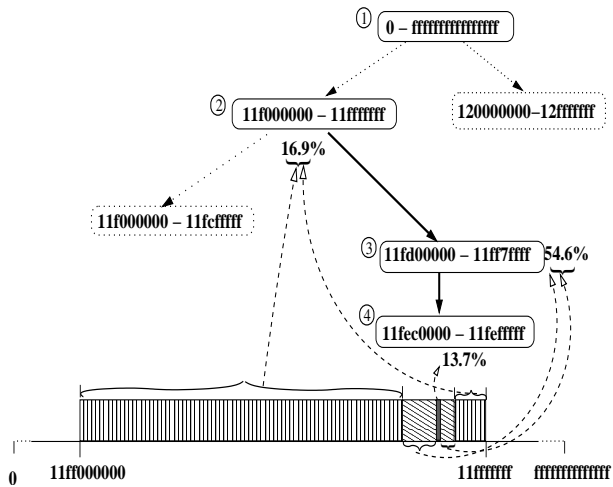


Figure 10: Memory-value profile characterized by RAP for `gcc` which identifies from which regions of the memory most zeros are being loaded. The horizontal axis represents the entire memory space. The hot nodes are labeled 1 through 4 and hot regions are shown as pattern-filled boxes

set of all load values which were subject to a cache miss we can quickly quantify this effect. Figure 9 shows the results of performing this analysis averaged over a set of benchmarks. The  $x$ -axis shows  $\log(\text{range\_width})$  of the different hot regions captured by RAP. The  $y$ -axis shows the coverage of all events, either loads, DL1 Cache misses, or DL2 Cache misses (depending on the curve). Take for example, DL1 misses. Hot-ranges (those ranges accounting for 10% or more of all DL1 misses) with a size of  $2^{16}$  or less account for about 56% of all DL1 misses. Looking at this figure, it is clear that in fact the value locality of cache misses is *more* than the value locality of all loads.

**Narrow Operand Profiling** – Another application of RAP would be finding regions of code with narrow operands. Find-

ing these regions might benefit operand width prediction and/or bit-width optimized compilation methods. We could build a RAP tree over the set of all instruction PCs which have a narrow operand (for example less than 16 bits). We profiled `gcc` and observed that the narrow-width operations were concentrated in very specific code regions, such as the file `flow.c` which accounted for 38.7% of all narrow-width operations. Within this file, the procedure `propagate_block` accounted for 31%, and a small block in this procedure which processed the live registers accounted for 6.4%.

**Zero-load Memory Ranges** – A different but related type of profile is to find out which regions of the data memory are responsible for load of a particular value, for example zero. This memory-value profiling could be used to guide bus compression schemes or track potentially inefficient data structures. Figure 10 shows a RAP tree for `gcc` built over the set of all memory addresses from which a zero was loaded. If the optimizers goal was to reduce the number of zero-loads, these memory ranges would be the best place to target. The horizontal axis represents the entire range of data memory (0-ffffffff). We focus on the hot nodes identified by RAP (labeled 1-4). We have zoomed in to show how RAP precisely identified distinct ranges which accounted for 16.9% (Node 2), 54.6% (Node 3) and 13.7% (Node 4) of the zero loads. For example, the address ranging from 11fd00000-11ff7fff (Node 3) accounts for a total of  $13.7\% + 54.6\% = 68.3\%$  of all zero loads in `gcc`. In fact, it was also observed that any load to this region has about 38% percent chance of being a zero.

In general, any event (cache misses, 0-loads, exceptions, ...) can be mapped using RAP, to the code that caused them, the memory address that was referred to, or the value on which an instruction operated. While the above profiling scenarios are not complete optimizations, they provide evidence that RAP has the potential to be both general purpose across many different types of profiles, and powerful enough to encourage new types of profiling.

## 5 Related Work

Range Adaptive Profiling is a novel method to provide hierarchical summary information on a stream of events. While we present a hardware based framework for dealing with vast amounts of profiling data, our technique builds on the profiling work of many other researchers. In this section we briefly summarize some of this work and relate it to our own contributions. We classify our related work into two broad categories:

**Software Based Profiling** - Software systems can be either statically instrumented with instrumentation tools such as ATOM [38] or dynamically through just-in-time compilers [24]. In software profiling, most of the effort has been spent on reducing the performance overhead of instrumentation such as through sampling [2] or bursty tracing [21]. Dynamic hot-path prediction techniques are described in [14]. Value profiles are another important form of profiles [6], which identify value invariance and proposes optimizations

through Convergent Profiling. There is software work on sampling more intelligently and even on compressing trace information to reduce the overheads involved. Larus [26] provides a technique to capture, in a compressed form, a program’s dynamic control flow. The idea of using software to extract a hierarchy of information using grammars, has been used to implement efficient data prefetching mechanisms [7, 8]. A general purpose software framework for dealing with compressed profile data is proposed in [43]. While these are powerful software mechanisms, they are not directly applicable to the problem of managing a very small number of hardware counters to enable high-throughput hardware-only profiling.

**Hardware Assisted Profiling** - The current industrial practice in hardware performance monitoring is performance counters, and several modern machines now support this idea [11, 12, 22]. These simple counter based schemes, while useful, suffer from a lack of flexibility and require significant software management in order to extract useful information [1]. Many researchers have examined the next steps that hardware assisted profiling should take. Proposed schemes range from those which use existing hardware on the processor to gather information which is later processed by a software program [1, 10, 32], to programmable profiling coprocessors [47]. [9] uses profile buffers to collect and analyze information. Sastry et. al. in [34] provide a framework for designing a variety of stream compressors and propose the stratified sampling scheme. An extension of the stratified sampling scheme is proposed by [31] which aims at reducing the cost of delivering gathered profile and proposes multi-hash and interval based profiling. Though these schemes provide efficient ways to process data they are not flexible enough to accommodate general queries. ProfileMe [13] and Relational Profiling Architecture are flexible and versatile schemes for gathering profile information. Zilles and Sohi [47] in their co-processor approach, design hardware to analyze the stream and compress it to provide concise and distilled profile information to the main processor. It has the ability to consider only a subset of the instructions for profiling and refocus resources after an instruction has been sufficiently characterized.

Our approach is orthogonal to most of the above approaches because RAP concentrates on building a useful on-line summary of the data, no matter what method is used to gather the data. RAP can be completely implemented in hardware and has the ability to efficiently identify the most important ranges of the profile and provide accurate information on the entire profile with very low overheads. We believe that there are important similarities between profiling a program executing billions of instructions per second and trying to monitor and analyze high speed networks [19, 15, 25] and that there is potential for further research along these lines. Indeed, RAP has been designed to be adaptable to a variety of different data streams that need to be processed at very high speed, and may even be applied in analyzing network traffic.

## 6 Conclusions

Amdahl’s law shows us that the common case is most important to performance so it makes sense to bias allocated resources towards the common case. The problem is that the common case changes as the program executes and we end up with a chicken-and-egg type of problem. In this paper we present Range Adaptive Profiling - a novel scheme to efficiently, adaptively and intelligently summarize high bandwidth streams of profile data. It allows users to specify a parameter ( $\epsilon$ ) which bounds the error with respect to the size of the input stream and also provides guarantees on worst case memory bounds independent of the size of the input stream *in a fully streaming fashion* (with only one-pass). This method can be applied to software profiling, and with the use of a specialized pipelined architecture, can be accelerated with hardware.

While it is not yet clear whether Range Adaptive Profiling will be general purpose enough to cover all profile types of interest, we have shown that it can make sense for summarizing at least three profile types: load values of cache misses, instruction PCs of narrow width operands, and memory addresses of zero-loads. The applicability of RAP can be further extended with multi-dimensional profiling which allows adaptive ranges over two or more variables. With this extension it is possible to handle edge profiles, data-code correlation studies, and general tuple space profiles, the details of which are beyond the scope of this paper. It may further be possible to unify our proposed techniques with existing sampling based schemes to create a single general purpose profiling system. While this future work may prove fruitful, to guide our initial algorithm and hardware design we have used load values and code profiling to measure the overheads of RAP and also to show the versatility of the scheme.

While some have shown the frequency of the top 50 individual loaded values in a program which might cover 40% of the program, our technique can *automatically generate range summaries* which include every value loaded in an entire SPEC benchmark, and we believe this type of analysis to be the first of its kind. This information could be used to drive many run-time optimizations including code specialization, value prediction, and bus encoding. While RAP has good worst case bounds, in the common case it is even better. For a set of benchmark programs from SPEC, we can provide 98% accurate information about hot code regions with only 8k bytes of memory and 99.73% accurate information with 64k bytes of memory. The RAP method is suitable for intelligent processing of the many different profile streams that may be generated from either a processor or computer network, and our future work will extend this technique to handle new forms of profiling in the processor.

## Acknowledgments

The authors would like to thank Cliff Young and the anonymous reviewers for their helpful feedback. This research was funded in part by National Science Foundation Grant CCF-0514738 and NSF Career Grant CCF-0448654.

## References

- [1] J. Anderson, W. Weihl, L. Berc, J. Dean, S. Ghemawat, M. Henzinger, S. Leung, R. Sites, M. Vandevoorde, and C. Waldspurger. Continuous Profiling: Where Have All the Cycles Gone? *ACM Transactions on Computer Systems (TOCS)*, 15(4):357–390, November 1997.
- [2] M. Arnold and B. Ryder. A Framework for Reducing the Cost of Instrumented Code. In *SIGPLAN Conference on Programming Language Design and Implementation*, pages 168–179, June 2001.
- [3] T. Ball and J. R. Larus. Efficient Path Profiling. In *International Symposium on Microarchitecture*, pages 46–57, Dec 1996.
- [4] D. Brooks and M. Martonosi. Dynamically Exploiting Narrow Width Operands to Improve Processor Power and Performance. In *Proceedings of the The Fifth International Symposium on High Performance Computer Architecture*, page 13, Washington, DC, USA, January 1999. IEEE Computer Society.
- [5] B. Buck and J. K. Hollingsworth. An API for Runtime Code Patching. *The International Journal of High Performance Computing Applications*, 14(4):317–329, 2000.
- [6] B. Calder, P. Feller, and A. Eustace. Value Profiling. In *International Symposium on Microarchitecture*, pages 259–269, December 1997.
- [7] T. Chlimbi. Efficient Representations and Abstractions for Quantifying and Exploiting Data Reference Locality. In *Conference on Programming Languages Design and Implementation*, June 2001.
- [8] T. Chlimbi and M. Hirzel. Dynamic Hot Data Stream Prefetching for General-Purpose Programs. In *Conference on Programming Languages Design and Implementation*, June 2002.
- [9] T. M. Conte, K. N. Menezes, and M. A. Hirsch. Accurate and Practical Profile-driven Compilation Using the Profile Buffer. In *International Symposium on Microarchitecture*, pages 36–45, December 1996.
- [10] T. M. Conte, B. A. Patel, and J. S. Cox. Using Branch Handling Hardware to Support Profile-driven Optimization. In *International symposium on Microarchitecture*, pages 12–21, November 1994.
- [11] Digital Equipment Corporation. Alpha 21164 Microprocessor Hardware Reference Manual. 1995.
- [12] Intel Corporation. Pentium(r) Pro Processor Developer’s Manual. In *McGraw-Hill*, June 1997.
- [13] J. Dean, J. Hicks, C. Waldspurger, W. Weihl, and G. Chrysos. ProfileMe: Hardware Support for Instruction-Level Profiling on Out-of-Order Processors. In *International Symposium on Microarchitecture*, pages 292–302, December 1997.
- [14] E. Duesterwald and V. Bala. Software Profiling for Hot Path Prediction: Less is More. In *Architectural support for programming languages and operating systems*, pages 202–211, October 2000.
- [15] C. Estan, S. Savage, and G. Varghese. Automatically Inferring Patterns of Resource Consumption in Network Traffic. In *SIGCOMM ’03: Proceedings of the 2003 conference on Applications, technologies, architectures, and protocols for computer communications*, Karlsruhe, Germany, August 2003.
- [16] First Silicon Solutions. Home page. <http://www.fs2.com>.
- [17] R. Gupta, D. Berson, and J.Z. Fang. Path profile guided partialredundancy elimination using speculation. In *IEEE International Conference on Computer Languages (ICCL)*, pages 230–239, May 1998.
- [18] T. Heil and J. E. Smith. Relational Profiling: Enabling Thread-level Parallelism in Virtual Machines. In *International symposium on Microarchitecture*, pages 281–290, 2000.
- [19] J. Hershberger, N. Shrivastava, S. Suri, and C. Toth. Adaptive Spatial Partitioning for Multidimensional Data Streams. In *International Symposium on Algorithms and Computation (ISAAC)*, 2004.
- [20] Hewlett-Packard. PA-RISC 1.1 Architecture and Instruction Set Reference Manual. 1994.
- [21] M. Hirzel and T. Chlimbi. Bursty Tracing: A Framework for Low-overhead Temporal Profiling. In *4th ACM Workshop on Feedback-Directed and Dynamic Optimization (FDDO-4)*, December 2001.
- [22] MIPS Technologies Inc. Mips R10000 Microprocessor User’s Manual. 1995.
- [23] Q. Jacobson, E. Rotenberg, and J.E. Smith. Path-based next trace prediction. In *30th IEEE/ACM International Symposium on Microarchitecture (MICRO)*, December 1997.
- [24] A. Krall. Efficient JavaVM Just-in-Time Compilation. In *International Conference on Parallel Architectures and Compilation Techniques*, pages 205–212, 1998.
- [25] C. Kruegel, F. Valeur, G. Vigna, and R.A. Kemmerer. Stateful Intrusion Detection for High-Speed Networks. In *IEEE Symposium on Security and Privacy*, pages 285–293, 2002.
- [26] J. Larus. Whole program paths. In *Conference on Programming Languages Design and Implementation (PLDI)*, pages 259–269, May 1999.
- [27] X. Li, Z. Liu, W. Li, and B. Liu. SCP-TCAM: A Power-Efficient Search Engine for fast IP Lookup. In *ISBN Proceedings*, 2004.
- [28] M. H. Lipasti and J. P. Shen. Exceeding the dataflow limit via value prediction. In *29th IEEE/ACM International Symposium on Microarchitecture (MICRO)*, pages 226–237, December 1996.
- [29] G. H. Loh. Exploiting data-width locality to increase superscalar execution bandwidth. In *MICRO 35: Proceedings of the 35th annual ACM/IEEE international symposium on Microarchitecture*, pages 395–405, Los Alamitos, CA, USA, 2002. IEEE Computer Society Press.
- [30] C. Luk and et. al. Pin: Building customized program analysis tools with dynamic instrumentation. In *Conference on Programming Languages Design and Implementation (PLDI)*, 2005.
- [31] S. Narayanasamy, T. Sherwood, S. Sair, B. Calder, and G. Varghese. Catching Accurate Profiles in Hardware. In *Int. Symp. on High-Performance Computer Architecture*, pages 269–280, February 2003.
- [32] R. V. Peri, S. Jinturkar, and L. Fajardo. A Novel Technique for Profiling Programs in Embedded Systems. In *ACM Workshop on Feedback-Directed and Dynamic Optimization*, 1999.
- [33] S. Rubin, R. Bodik, and T. Chlimbi. An efficient profile-analysis framework for data layout optimizations. In *The 29th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*, Jan 2002.
- [34] S. Sastry, R. Bodik, and J. Smith. Rapid Profiling via Stratified Sampling. In *Annual International Symposium on Computer Architecture*, pages 278–289, July 2001.
- [35] P. Shivakumar and N. P. Jouppi. Cacti 3.0: An Integrated Cache Timing, Power and Area Model, 2001/2.
- [36] V. Srinivasan and G. Varghese. Fast Address Lookups Using Controlled Prefix Expansion. *ACM Transactions on Computer Systems*, 7(1):1–40, February 1999.
- [37] A. Srivastava, A. Edwards, and H. Vo. Vulcan: Binary transformation in a distributed environment. technical report msr-tr-2001-50, microsoft research. April 2001.
- [38] A. Srivastava and A. Eustace. ATOM: A System for Building Customized Program Analysis Tools. In *Conference on Programming Languages Design and Implementation*, pages 196–205, 1994.
- [39] H. Wang, X. Zhu, L. Peh, and S. Malik. Orion: A Power-Performance Simulator for Interconnection Networks. In *Proceedings of the 35th International Symposium on Microarchitecture*, Istanbul, Turkey, November 2002.
- [40] J. Yang and R. Gupta. Frequent Value Locality and its Applications. In *ACM Transactions on Embedded Computing Systems*, 2002.
- [41] J. Yang, Y. Zhang, and R. Gupta. Frequent value compression in data caches. In *International Symposium on Microarchitecture*, pages 258–265, 2000.
- [42] C. Young and M.D. Smith. Better global scheduling using path profiles. In *IEEE/ACM International Symposium on Microarchitecture (MICRO)*, pages 115–123, 1998.
- [43] X. Zhang and R. Gupta. Whole execution traces. In *Proceedings of the 37th International Symposium on Microarchitecture*, pages 105–116, Washington, DC, USA, 2004. IEEE Computer Society.
- [44] Y. Zhang and R. Gupta. Timestamped whole program path representation and its applications. In *Conference on Programming Languages Design and Implementation (PLDI)*, pages 180–190, June 2001.
- [45] H. Zhou, J. Flanagan, and T. M. Conte. Detecting global stride locality in value streams. In *Proceedings of the 30th Annual International Symposium on Computer Architecture*, pages 324–335, New York, NY, USA, 2003. ACM Press.
- [46] P. Zhou, F. Qin, W. Liu, Y. Zhou, and J. Torrellas. iWatcher: Efficient Architectural Support for Software Debugging. In *Proceedings of the 31st annual International Symposium on Computer Architecture (ISCA’04)*, June 2004.
- [47] C. Zilles and G. Sohi. A Programmable Co-Processor for Profiling. In *The Seventh International Symposium on High Performance Computer Architecture*, January 2001.