Estimators Also Need Shared Values to Grow Together

Erez Tsidon*§, Iddo Hanniel* and Isaac Keslassy*
* Technion § Qualcomm
{erezts@tx., ihanniel@, isaac@ee.}technion.ac.il

***Abstract*—Network management applications require large numbers of counters in order to collect traffic characteristics for each network flow. However, these counters often barely fit into on-chip SRAM memories. Past papers have proposed using counter estimators instead, thus trading off counter precision for a lower number of bits. But these estimators do not achieve optimal estimation error, and cannot always scale to arbitrary counter values.**

In this paper, we introduce the CEDAR algorithm for decoupling the counter estimators from their estimation values, which are quantized into estimation levels and shared among many estimators. These decoupled and shared estimation values enable us to easily adjust them without needing to go through all the counters. We demonstrate how our CEDAR scheme achieves the min-max relative error, i.e., can guarantee the best possible relative error over the entire counter scale. We also explain how to use dynamic adaptive estimation values in order to support counter up-scaling and adjust the estimation error depending on the current maximal counter.

Finally we implement CEDAR on FPGA and explain how it can run at line rate. We further analyze its performance and size requirements.

I. Introduction

A. Background

In many high-end networking devices, the amount of *on-chip memory needed by counters* keeps growing and starts reaching its limits. This is because the needed number of counters in networking devices keeps increasing with the number of flows. It also keeps growing with the number of measurement-based applications, since each application can require different types of counters per flow. In addition, not only the number of counters, but also the counter sizes keep increasing with the line rates, since line rates directly determine the maximum counter value.

This memory limit on the number of counters can directly affect many algorithms for networking devices, since these often rely on the assumption that packet counts and flow rates are easily available. For instance, the ability to easily measure and compare flows lies at the basis of many fairness, scheduling, flow control, admission control, load-balancing, shaping and policing algorithms.

As a simplified example, consider a current high-speed multi-stage router [1]–[3], e.g. with 256 input and output ports and 4 priority classes. Then a middle-stage element sees $256^2 \cdot 4 \approx 262,000$ unicast router flows. Keeping only three counters per flow (e.g., the number of packets of the flow in the switch element, and its current arrival and departure rates) with 16 bits per counter would yield a needed memory size of 48 bits per flow, i.e. a total of 12 Mb $= 1.5$ MB. After implementing all the other functions in the middle-stage element, such as packet switching and scheduling control, this can fit into available commodity on-chip SRAM memory (barely). However, in next-generation routers with $1,024$ ports, the needed memory would scale to 24 MB—well beyond commodity SRAM sizes [4].

In this paper, we consider networking device applications for which counter *estimates* are sufficient, even if the estimates have small errors, as long as all those estimates can fit on commodity on-chip SRAM memory and can be accessed in real time. For instance, a flow control mechanism similar to AF-QCN [5] could use rate estimates for each flow to obtain quick convergence to fairness. Taken to the extreme, these algorithms might in fact only depend on the order of magnitude of the estimated flow rate. For example, a differentiated service algorithm might provide a different service to flows with estimated rates of about 10Kbps, 100Kbps or 1Mbps.

This paper is about enabling many per-flow counters in a scalable way, both for packet counts and packet rates. We do so by replacing *exact counters* with *shared counter estimates*. These estimates will enable us to trade off a decreased needed amount of memory with an increased error probability in the counter estimation. In addition, using the estimator sharing, we provide ways to smoothly *change the counter scale*, and therefore adapt the estimation error to the traffic size in a dynamic way. This is particularly useful for estimates of window-based average *rates*, since they rely on counters that keep increasing and therefore need to adapt over several orders of magnitude.

B. Alternative Solutions and Related Work

The following alternative solutions attempt to use less bits per counter.

A first approach is to *keep less counters*. A straightforward solution adopted by high-end device designers against the memory size problem is to avoid keeping an on-chip counter for each flow, and instead keep a counter *per pre-determined logical flow aggregate*. For instance, keep a counter for all packets with a certain priority, or for all packets coming from a given input port. However, this might prevent the use of several algorithms. For instance, upon congestion, this prevents the use of fairness algorithms based on the flow rates, such as the flow control and differentiated-service algorithms previously

mentioned. Another way to keep less counters is to only keep counters for heavy hitters [6], [7]. However, while this can help in many applications, our goal is to provide a counter estimate for *all* flows.

A second approach is to only *keep counter estimates for active flows* using a structure based on *CBF* (Counting Bloom Filter) [8]. For instance, [9] suggests the Conservative Update algorithm to reduce errors. Likewise, [10], [11] introduce the Counter Braids structure to achieve near-entropy compression. However, the Counter Braids algorithms are not designed for *online* counter estimation, which is assumed in this paper. Also, both of these structures perform most efficiently when most counter values are zero, which is not necessarily the case. In practice, when most counter values are zero, CBF can be combined with our estimators to provide even better error rates for the same amount of memory: instead of estimating flow counters, our estimators would simply *estimate the CBF entry counters*.

Likewise, a related algorithm, BRICK [12], restricts the counter sum in order to compress counters. This requirement assumes a flow distribution that keeps most of the counter values at zero. In this paper we attempt to avoid such an assumption.

A third approach is to use *DRAM*, either by combining small on-chip SRAM counters with large off-chip DRAM counters [13]–[16], or by only using off-chip DRAM counters with potentially a small queue in SRAM [17]–[19]. This DRAM-based approach provides exact counters instead of counter estimates. It can also obtain cheaper solutions by relying on DRAM rather than SRAM. Unfortunately, because it relies on off-chip DRAMs with longer access latencies, this approach cannot satisfy arbitrary real-time counter reads at line rate. For instance, it may need to restrict the pattern of the counter queries and/or add a lookahead queue of unpredictable length. In addition, this approach forces the designer to deal with the buffer management policy for the off-line DRAM in addition to the SRAM that is used in other applications, which may be a barrier to its use.

A fourth approach is to use *counter approximations* instead of real values, as in the Approximate-Counting [22], SAC [20] and DISCO [21] algorithms. In these methods each counter is replaced with an estimator that requires significantly less bits (e.g. 12 bits vs. 32). However, each method restricts its estimation value distribution to predefined values. For example, SAC restricts its estimation values to be of the form $a2^b$. We will show that these estimation values are not necessarily optimal and provide optimal values. In addition, some of these methods (e.g. DISCO) do not support up-scaling procedures, thus the maximal possible estimated value must be known ahead of time, resulting in a statically large error over the entire counter scale. We propose a method to adjust the estimation error to the current maximal counter value.

C. Our Contributions

In this paper we introduce a counter architecture, called CEDAR (Counter Estimation Decoupling for Approximate Rates), which decouples counter values from flow symbols. The goal of CEDAR is to provide *real-time scalable counter estimates*. CEDAR gathers estimation values into one array with a relatively small number of entries. Flow symbols are used as pointers into this estimation array. Since the number of estimation array entries is low compared to the number of flow entries (e.g. hundreds vs. millions), we show that the additional required memory is negligible.

Using this estimator array, we have the freedom to choose any estimation values. We exploit this advantage by determining the exact estimators that minimize the maximal relative error over the entire counter scale. That is, for any counter value, CEDAR can guarantee that the relative error does not exceed a predetermined guaranteed value, which is proved to be minimal. This estimation is particularly useful for applications that require fast reading and writing of flow counters with a predetermined relative estimation error that is independent of the counter value.

Then, to obtain counter estimators without a predetermined maximum, we introduce an online up-scaling scheme. This scheme is useful when there is no advance knowledge about the flow distribution, and therefore the maximal counter value cannot be predicted. Our up-scaling algorithm adjusts the relative error of the entire counter scale to the current maximal counter value. In this way the estimation error at any time is uniquely derived from the stream behavior.

In order to evaluate our algorithms, we use Internet traces and compare the estimation error of CEDAR versus SAC [20] and DISCO [21]. We show how our algorithm can more efficiently estimate counter values given the same amount of memory. We also show how our up-scaling mechanism keeps the max-relative-error adjusted to the counter scale.

Finally, we demonstrate the CEDAR implementation on FPGA and analyze its performance and size requirements.

Paper Organization: The CEDAR architecture is defined in Section II. Then Section III explains how to attain the min-max-relative-error. Next, Section IV describes the CEDAR up-scaling algorithm. In Section V we evaluate the CEDAR performance on real Internet traces. Last, we introduce the FPGA implementation in Section VI.

II. The CEDAR Architecture

A. Architecture

As shown in Figure 1, the CEDAR (Counter Estimation Decoupling for Approximate Rates) architecture is based on two arrays. First, a flow array F of size N that contains pointers into the estimator array A for each of the flows F_j, $0 \leq j \leq N-1$. Second, an estimation array A of size L that contains estimator values A_i, $0 \leq i \leq L-1$.

For instance, in the figure, the first flow F_0 points to A_1 (the pointer to index 1 is denoted $p(1)$), and the value of estimator A_2 is 1.2. Therefore the estimator value for flow F_0 is 1.2.

Assuming that L is a power of 2, the number of bits needed for each of the pointers in F is simply $\log_2 L$. We also assume that each estimator value in A uses q bits.

B. Conversion to the CEDAR Structure

We show now that *any counter estimation method* that represents each estimate independently using a fixed number

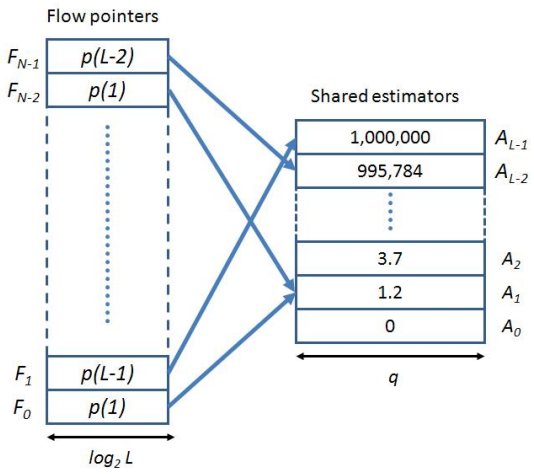

Fig. 1. CEDAR structure. The flow pointers on the left point to shared estimators on the right. For example, the estimator for flow F_0 is $A_1 = 1.2$

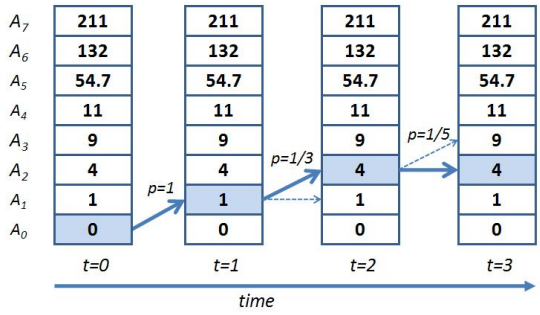

Fig. 2. Update mechanism of the CEDAR algorithm given consecutive packet arrivals for a given flow. The flow counter estimate is initialized at $A_0 = 0$. After the first arrival, the difference between A_0 and A_1 is 1, therefore the update step is deterministic with probability $p = 1/1 = 1$ and the new estimate is $A_1 = 1$. After the second arrival, $A_2 - A_1 = 3$, therefore the flow pointer is incremented with probability $1/3$ to $A_2 = 4$. At the third arrival, $A_3 - A_2 = 5$, therefore the increment probability is $1/5$. Here it is not incremented.

of bits can be converted easily into a CEDAR structure with asymptotically negligible effect on the estimation error. This applies for instance to the SAC [20] and DISCO [21] counter estimation methods.

Theorem 1: Any counter estimation method using a fixed number of bits per counter estimate can be converted into the CEDAR structure with the same estimation error, and an asymptotically negligible overhead as $N \to \infty$.

Proof: Consider an estimation method that uses q bits per counter for counter estimation. Thus, it can have a maximum of 2^q distinct symbols. If we build from these symbols an estimation array of size $L = 2^q$, then we can use the counter symbols as pointers into this estimator array and run the original algorithm without any change. Therefore we get a CEDAR structure that keeps the same estimation error. In addition the structure overhead is the additional contribution of the estimator array, i.e. $\frac{q \cdot 2^q}{N \cdot q} = \frac{2^q}{N} \xrightarrow{N \to \infty} 0$. ■

C. Performance Measures

In the paper we only use unit increments that represent packet arrivals. Note that the estimation error with variable increments (e.g. when measuring bytes) is actually lower given the same scale, because it enables more precise scale jumps, and therefore using unit increments results in a worst-case analysis.

We denote by $T(l)$ the random variable that represents the required amount of traffic for a certain flow to point to estimator A_l, i.e. its *hitting time*. As in [20], [21], assuming our algorithms yield unbiased estimators and assuming $\mathbb{E}[T(l)] = A_l$, we will define the *relative error* (or coefficient of variation) as

$$\delta[T(l)] = \sqrt{\frac{\mathrm{Var}[T(l)]}{(\mathbb{E}[T(l)])^2}} = \frac{\sigma[T(l)]}{\mathbb{E}[T(l)]}$$

D. Static CEDAR Algorithm

We start by introducing the static version of CEDAR, which uses a static and pre-determined estimator array A. Later, in Section IV, we will introduce the dynamic version of CEDAR that enables scaling the estimator array.

In the static CEDAR algorithm, we make two simplifying assumptions on the fixed CEDAR structure. First, we assume that the estimator values A_i are defined in ascending order, with $A_0 = 0$. We denote the differences between two successive estimators as $D_i = A_{i+1} - A_i$, where $0 \le i \le L-2$, and further assume that $D_i \ge 1$ for any i. For instance, in Figure 1, $D_0 = 1.2 - 0 = 1.2$ and $D_1 = 3.7 - 1.2 = 2.5$.

As The static CEDAR algorithm works as follows. At the start, all flow counters point to $A_0 = 0$. Then, at each arrival of a packet from flow j, which corresponds to a unit increment of its real counter value, the flow's pointer F_j is incremented by one with probability $\frac{1}{D_{F_j}} = \frac{1}{A_{F_j+1} - A_{F_j}}$. In other words, if the flow pointer first points to estimated value A_i, then it either starts pointing to A_{i+1} with probability $1/D_{F_j}$, or continues pointing to A_i with probability $1 - 1/D_{F_j}$. Note that since $D_i \ge 1$ for any i, these probabilities are well defined.

Figure 2 further illustrates an example of such update sequence in the static CEDAR algorithm.

E. Unbiasedness and Traffic Expectation

Given our definition of the static CEDAR algorithm, we first show that the expectation of random variable $T(l)$ equals the estimated value A_l. This means that the expected hitting time of A_l is exactly A_l, i.e. the expected number of arrivals needed to reach some estimated value is exactly this estimated value.

Theorem 2: The hitting time $T(l)$ of estimator A_l satisfies the following property: $\mathbb{E}[T(l)] = A_l$

Proof: We can divide $T(l)$ into l i.i.d geometric random variables $G(p_i)$ with parameter $p_i = \frac{1}{D_i}, i = [0, ..., l-1]$, therefore

$$\mathbb{E}[T(l)] = \mathbb{E}\left[\sum_{i=0}^{l-1} G(\frac{1}{D_i})\right] = \sum_{i=0}^{l-1} \mathbb{E}\left[G(\frac{1}{D_i})\right] = \sum_{i=0}^{l-1} D_i = A_l$$

■

In addition, regardless of the exact values used within estimation array A, the CEDAR estimation is unbiased, as long as we assume that the estimation scale is unbounded (or

that its maximum is never reached). The unbiasedness proof of CEDAR is similar to the proof of Theorem 1 in DISCO [21], therefore we do not repeat it.

Property 1: CEDAR estimators are unbiased.

III. Relative Error Minimization

A. Min-Max Relative Error

A network device designer will want to obtain a guarantee that *the counter estimates have good performance no matter how large the counters are*. For instance, such a guarantee might be that the relative error is below 5% over the whole scale of counters.

Unlike former approaches that restrict the counter estimators to specific values, the flexibility of the CEDAR structure enables us to analytically determine and implement the *optimal estimation array* that achieves such a guarantee. In the remainder of this section, we first determine the values of this optimal estimation array in Theorem 3, then prove that it is indeed optimal in Theorem 4.

More specifically, given an estimation array size L and a fixed maximal estimator value A_{L-1}, we want to determine the set of estimator values $\{A_0, ..., A_{L-1}\}$ that minimizes the maximum relative error over all estimators, i.e. achieves $\min_A \max_l \delta[T(l)]$. In other words, this set should minimize the guaranteed error δ such that for all ℓ, $\delta[T(l)] \leq \delta$.

As we saw in Theorem 2, $T(l)$ is a sum of l i.i.d geometric random variables, therefore

$$\delta[T(l)] = \sqrt{\frac{\mathrm{Var}[T(l)]}{(E[T(l)])^2}} = \sqrt{\frac{\sum_{i=0}^{l-1} \frac{1-1/D_i}{1/D_i^2}}{\left(\sum_{i=0}^{l-1} D_i\right)^2}}$$

and we can rewrite $\delta[T(l)] \leq \delta$ for all ℓ as

$$\forall l : F_l(D_0, ..., D_l) = \sum_{i=0}^{l} \left(D_i^2 - D_i\right) - \delta^2 \left(\sum_{i=0}^{l} D_i\right)^2 \leq 0 \tag{1}$$

We later show in Theorem 4 that in order to achieve such min-max relative error, all estimators must reach an equal relative error δ, that is,

$$\forall l : \delta[T(l)] = \delta,$$

which we can rewrite as $\forall l : F_l(D_0, ..., D_l) = 0$. We first prove in the following theorem that we achieve such an *equal-relative-error* estimation array iff the estimation values are given by the following recursive equation,

$$D_l = \frac{1 + 2\delta^2 \sum_{i=0}^{l-1} D_i}{1 - \delta^2}, \tag{2}$$

with $D_0 = \frac{1}{1-\delta^2}$, before proving in Theorem 4 that we indeed need to achieve this equal-relative-error to provide an optimal guarantee.

Theorem 3: Given a target relative error δ, and the $L-1$ constraints: $\forall l \in [0, L-2], F_l = 0$, then the estimation value increments D_l are necessarily given by the recursive function in Equation (2).

Proof: First, we highlight the recursive nature of the F_l functions. Extracting D_l from F_l yields:

$$\begin{aligned} F_l =& \sum_{i=0}^{l-1} D_i^2 + D_l^2 - \sum_{i=0}^{l-1} D_i - D_l \\ & - \left((\sum_{i=0}^{l-1} D_i)^2 + 2D_l \sum_{i=0}^{l-1} D_i + D_l^2 \right) \delta^2 \\ =& (1-\delta^2) D_l^2 - (2\delta^2 \sum_{i=0}^{l-1} D_i + 1) D_l \\ & + \underbrace{\sum_{i=0}^{l-1} D_i^2 - \sum_{i=0}^{l-1} D_i - (\sum_{i=0}^{l-1} D_i^2)^2 \delta^2}_{F_{l-1}} = 0 \end{aligned} \tag{3}$$

Assigning $F_{l-1} = 0$ immediately gives Equation (2). ∎

The following theorem proves that minimizing the maximal relative error δ for a given A_{L-1} is equivalent to maximizing A_{L-1} for a given maximal relative error δ. More significantly, it demonstrates that both are equivalent to equalizing all relative errors, therefore yielding the recursive equation obtained above (Theorem 3).

Theorem 4: Given the $L-1$ constraints: $\forall l \in [0, L-2], F_l \leq 0$, the following claims are equivalent:

(i) $A_{L-1} = \sum_{i=0}^{L-2} D_i$ is maximal *(maximizing A_{L-1})*.
(ii) $\forall l \in [0, L-2], F_l = 0$ *(equalizing all relative errors)*.
(iii) δ is the minimal guaranteed relative error *(minimizing δ)*.

Proof: We start by proving that (i)⇒(ii). Assuming $\sum_{i=0}^{L-2} D_i$ is maximal, the $(L-1)^{th}$ equality, $F_{L-2} = 0$, is given by the following lemma:

Lemma 1: Assume that $\forall l \in [0, L-2], F_l \leq 0$. If $A_{L-1} = \sum_{i=0}^{L-2} D_i$ is maximal then $F_{L-2} = 0$

Proof: We assume by contradiction that there is a set $\{D_0, ..., D_{L-2}\}$ that satisfies all $L-1$ constraints and achieves maximal $\sum_{i=0}^{L-2} D_i$ but $F_{L-2} < 0$. We choose $\epsilon > 0$ such that

$$\epsilon < \min\left(1, \frac{-F_{L-2} \sum_{i=0}^{L-2} D_i^2}{2D_{L-2}}\right)$$

We increase D_{L-2} by ϵ and thus increase the maximal sum $\sum_{i=0}^{L-2} D_i$ by ϵ. Changing D_{L-2} has no effect on the other $L-2$ constraints $F_l \leq 0, l = 0, ..., L-3$. Consider now the $(L-1)^{th}$ constraint:

$$\begin{aligned} & F_{L-2}(D_0, ..., D_{L-2} + \epsilon) = \\ & \frac{\sum_{i=0}^{L-3} D_i^2 + (D_{L-2} + \epsilon)^2 - \sum_{i=0}^{L-2} D_i - \epsilon}{(\sum_{i=0}^{L-2} D_i + \epsilon)^2} < \\ & \frac{\sum_{i=0}^{L-2} D_i^2 - \sum_{i=0}^{L-2} D_i}{(\sum_{i=0}^{L-2} D_i)^2} + \frac{\epsilon^2 - \epsilon + 2\epsilon D_{L-2}}{(\sum_{i=0}^{L-2} D_i)^2} < \\ & F_{L-2}(D_0, ..., D_{L-2}) + \frac{2\epsilon D_{L-2}}{(\sum_{i=0}^{L-2} D_i)^2} < 0 \end{aligned} \tag{4}$$

∎

We now prove that $F_{L-2} = 0 \Rightarrow F_{L-3} = 0$, and more generally, by induction, $F_l = 0 \Rightarrow F_{l-1} = 0$, i.e. $F_l = 0$ for all l. We assume by contradiction that $F_{L-3} < 0$ and show that this leads to $D_i = D_j$ for all i, j, but assigning identical

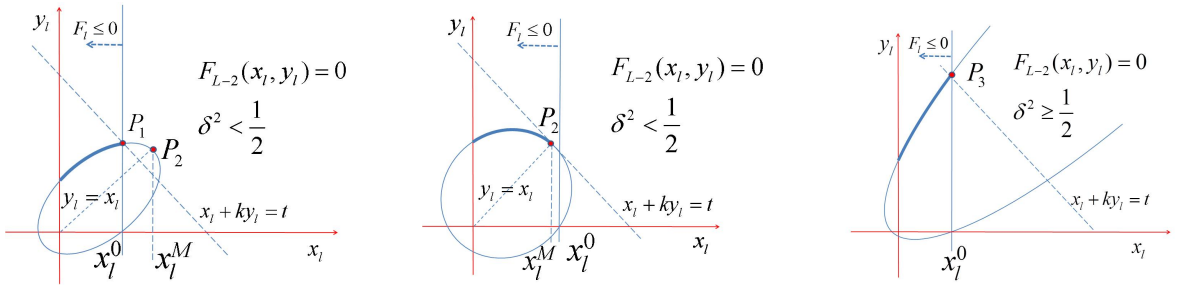

(a) Ellipse case, option 1: intersects with $F_l \le 0$ constraint to the left of the tangency point P_2. The maximum for $x_l + ky_l$ is attained at the intersection point P_1.

(b) Ellipse case, option 2: intersects with $F_l \le 0$ constraint to the right of the tangency point P_2. The maximum for $x_l + ky_l$ is therefore attained at P_2.

(c) Parabola / Hyperbola case: always intersect with $F_l \le 0$ constraint at the upper branch of the curve. The maximum for $x_l + ky_l$ is attained at the intersection point P_3.

Fig. 3. Geometric view of Theorem 4 based on quadratic curves: the quadratic curves represents the $F_{L-2} = 0$ constraint (according to the δ value) while the vertical line $x_l = x_l^0$ represents the $F_l \le 0$ constraint.

values does not result in a maximal $\sum_{i=0}^{L-2} D_i$ and therefore the assumption is refuted. We begin by assigning $l = L-3$ and follow the following steps in the proof:

1) We show that if $F_l < 0$ then the only way to achieve maximal $\sum_{i=0}^{L-2} D_i$ is when $F_{l-1} < 0$ and $D_l = D_{l+1}$.
2) For $l > 0$, we set $l \leftarrow l-1$ and return to 1)
3) When $l = 0$ we get $\forall i,j : D_i = D_j$. We will refute this result by comparing it to the values we got in Equation (2), which achieve a larger $\sum_{i=0}^{L-2} D_i$.

We wish to provide an intuition for the first recursion cycle where $l = L-3$, before proving all recursion cycles. Assuming $\sum_{i=0}^{L-2} D_i$ is maximal, Lemma 1 gives us the $(L-1)^{th}$ equality: $F_{L-2} = 0$. In a similar manner to the way we extracted D_{l-1} in Equation (3), we extract both D_{L-2} and D_{L-3} from F_{L-2}, resulting in:

$$\begin{aligned} F_{L-2} &= (1-\delta^2)D_{L-3}^2 - 2\delta^2 D_{L-2}D_{L-3} \\ &+ (1-\delta^2)D_{L-2}^2 - (2\delta^2 \sum_{i=0}^{L-4} D_i + 1)D_{L-3} \\ &- (2\delta^2 \sum_{i=0}^{L-4} D_i + 1)D_{L-2} + F_{L-4} = 0 \end{aligned} \quad (5)$$

Interestingly, we can use a *geometric approach* to this optimization. Let us set $l = L-3$ and denote the following symbols: $x_l = D_l; y_l = D_{l+1}; A = 1-\delta^2; B = -2\delta^2; C_1 = -(1+2\delta^2 \sum_{i=0}^{l-1} D_i)$. We can rewrite Equation (5) as:

$$Ax_l^2 + Bx_l y_l + Ay_l^2 + C_1 x_l + C_1 y_l + F_{l-1} = 0 \quad (6)$$

where $A > 0, B < 0, C_1 < 0$, and $F_{l-1} \le 0$. *This equation defines a rotated quadratic curve in the (x_l, y_l) plane:* an *ellipse* when $\delta < \sqrt{0.5}$, a *parabola* when $\delta = \sqrt{0.5}$, and a *hyperbola* when $\delta > \sqrt{0.5}$ (see Figure 3)[1]. We now consider the constraint $F_l \le 0$ (i.e., $F_{L-3} \le 0$) in the (x_l, y_l) plane using our new notations. Note that from Equation (3), we can see that F_{L-3} can be computed from F_{L-2} by assigning $D_{L-2} = 0$. In our new notation this means setting $y_l = 0$. We get $Ax_l^2 + C_1 x_l + F_{l-1} \le 0$, implying that

$$\frac{-C_1 - \sqrt{C_1^2 - 4AF_{l-1}}}{2A} \le x_l \le \frac{-C_1 + \sqrt{C_1^2 - 4AF_{l-1}}}{2A} \quad (7)$$

Since $x_l = D_l \ge 1$ and the left-hand side of Equation (7) is always negative, we are only left with the following constraint:

$$x_l \le x_l^0 = \frac{-C_1 + \sqrt{C_1^2 - 4AF_{l-1}}}{2A}$$

That is, as shown in Figure 3, the solution range is on the quadratic curve from the left side of the vertical line $x_l = x_l^0$. In the hyperbola/parabola case the vertical line always intersects with the upper branch of the quadratic curve at P_3 since the curve is unbounded in the x_l axis. Therefore, $x_l + ky_l = t$ gets its maximal value at the intersection point P_3 (in the specific case we are analyzing, of the first recursion cycle, $k = 1$). Notice that in the hyperbolic case the left asymptote cannot be vertical (which would prevent the curve and the vertical line from intersecting) since setting $x_l = 0$ at Equation (6) yields an intersection point between the hyperbola and the y axis.

In the ellipse case there are two options. Let $x_l = x_l^M$ be the vertical that passes through P_2, the point of the ellipse that is tangent to the line $x_l + ky_l = t$ (see Figure 3). In option 1 $x_l^0 < x_l^M$ and the maximal value is attained at P_1, whereas in option 2 $x_l^M < x_l^0$ and the maximal value is attained at the tangency point P_2. Later on, in Lemma 2, we show that P_2 is also the intersection point between the ellipse and $y_l = x_l$.

Since in constructing the quadratic curve we did not make any specific assumptions on the values of $D_i, i = 0, ..., L-3$ (except that $F_{L-4} \le 0$) then Figure 3 is in particular true for the values that attain the maximal $\sum_{i=0}^{L-2} D_i$ under the constraints. Note also that at points P_1 and P_3, $F_l = 0$. Therefore, when $\delta^2 \ge 0.5$ (the parabola/hyperbola case) there is always an intersection point at the upper branch of the quadratic curve and the contradiction assumption is refuted since the maximum is attained at $F_l = 0$. In the case of $\delta^2 < 0.5$, according to the contradiction assumption, $F_l < 0$, i.e. the vertical line cannot intersect the ellipse at P_1. Therefore,

[1]On how to classify a quadratic curve, see for example [23] or any basic textbook on curves and surfaces such as [24].

in this case we are in option 2 and the maximum is attained on the ellipse tangent at P_2 where $x_l = y_l$, i.e., when $D_l = D_{l+1}$.

The following lemma proves this formally for all recursion cycles:

Lemma 2: Under the constraint:

$$F = (1-\delta^2)x^2 - (1+2\delta^2\sum_{i=0}^{L-3-k} D_i)x - (2k\delta^2)xy + k(1-k\delta^2)y^2 - k(1+2\delta^2\sum_{i=0}^{L-3-k} D_i)y + F_{L-3-k} = 0 \quad (8)$$

(k is a positive integer, $F_{L-3-k} \leq 0$). The expression $x+ky$ is either unbounded or attains its maximum when $x = y$.

Proof: We can see that $F(x,y)$ is a continuously differentiable function and $x + ky$ is affine (On KKT's sufficient conditions see [25]). Therefore we can use Lagrange multipliers to find the extremum of $x + ky$ under the constraint. Note that $\partial(x + ky)/\partial x = 1$ and $\partial(x + ky)/\partial y = k$.

$$\frac{\partial F}{\partial x} = -2k\delta^2 y + 2(1-\delta^2)x - (1+2\delta^2\sum_{i=0}^{L-3-k} D_i) = \frac{1}{\lambda}$$

$$\frac{\partial F}{\partial y} = 2k(1-k\delta^2)y - 2k\delta^2 x - k(1+2\delta^2\sum_{i=0}^{L-3-k} D_i) = \frac{k}{\lambda}$$

eliminating λ by equating $\frac{\partial F}{\partial y} = k\frac{\partial F}{\partial x}$, we get:

$$2kx = 2ky.$$

That is, the extremum M of $x + ky$ under the constraint F is achieved when $x = y$.

Furthermore, let us denote by x_M the value of x (which is also the value of y) where M is attained, the derivative $d(x + ky)/dx$ is positive in the domain $0 \leq x \leq x_M$. This is easily shown by the fact that

$$dy/dx = -\frac{\partial F_{L-2}/\partial x}{\partial F_{L-2}/\partial y} = -1/k$$

at $x = x_M$. At $x = 0$, $dy/dx > 0$ (assign $x = 0$ and $y > 0$ at the partial derivative equations) and therefore, since F is convex,$d(x + ky)/dx \geq 1 + k(-1/k) = 0$ when $0 \leq x \leq x_M$. That is, $(x + ky)$ is an increasing function at the domain $0 \leq x \leq x_M$, attaining its maximum at x_M.

Thus, if $x + ky$ under the constraint F has a maximum M (e.g., when the quadratic curve F is an ellipse), M is attained when $x = y$. In case the quadratic curve F is unbounded (a hyperbola or parabola) then $x + ky$ does not have a maximum under the constraint F. ∎

Note that the private case of the lemma when $k = 1$ is the case we just explored in Equation (5), where $F_{L-2} = 0$ is the constraint. For this case, the maximum induced by the constraint is (by assigning $x = y$ and solving Equation (6)):

$$x^M_{L-3} = \frac{-2C_1 + \sqrt{4C_1^2 - 4(2A+B)F_{L-4}}}{2(2A+B)} \quad (9)$$

Now, if $x^M_{L-3} \geq x^0_{L-3}$ (option 1), then the maximum of $x + y$ (and therefore the maximum of $\sum_{i=0}^{L-2} D_i$) is attained at x^0_{L-3}, i.e., when $F_{L-3} = 0$. Also note that if $F_{L-4} = 0$, then $x^M_{L-3} \geq x^0_{L-3}$ since $-C_1/A < -2C_1/(2A+B)$.

By our contradiction assumption, $F_{L-3} < 0$, the maximum must be attained when $x^M_{L-3} < x^0_{L-3}$ and $F_{L-4} < 0$. Therefore, $F_{L-3} < 0 \Rightarrow F_{L-4} < 0$. By Lemma 2, the maximum must be attained when $x = y$, i.e., when $D_{L-3} = D_{L-2}$. This proves item (1) of the recursion steps for the first cycle.

Now, we can continue to the next recursion step $l = L - 4$. We can plug-in $D_{L-3} = D_{L-2} = D$ into $F_{L-2} = 0$ and get the following constraint in $x_l = D_{L-4}$ and $y_l = D_{L-3} = D_{L-2} = D$:

$$F = (1-\delta^2)x_l^2 - (1+2\delta^2\sum_{i=0}^{L-5} D_i)x_l - (4\delta^2)x_l y_l + 2(1-2\delta^2)y_l^2 - 2(1+2\delta^2\sum_{i=0}^{L-5} D_i)y_l + F_{L-5} = 0, \quad (10)$$

which is a private case for $k = 2$ of Lemma 2.

Applying the lemma onto the new constraint, we get that the maximal value of $x_l + 2y_l$ on F is attained when $x_l = y_l$.

In a similar manner to what was done above, we can formulate x^0_{L-4}, which is achieved when $F_{L-4} = 0$ and x^M_{L-4}, which is achieved when $x_l = y_l$ (i.e., when $D_{L-4} = D_{L-3} = D_{L-2} = D$). Again, if $x^M_{L-4} \geq x^0_{L-4}$, then the maximum of $x_l + 2y_l$ (and therefore the maximum of $\sum_{i=0}^{L-2} D_i$) is attained at x^0_{L-4}, i.e., when $F_{L-4} = 0$. But, this will mean that (in the level above) $x^M_{L-3} \geq x^0_{L-3}$, which violates our contradiction assumption.

Thus, the maximum must be attained when $x^M_{L-4} < x^0_{L-4}$ and $F_{L-5} < 0$. By the lemma, it is therefore attained when $x_l = y_l$, i.e., when $D_{L-4} = D_{L-3} = D_{L-2} = D$.

Applying the same logic recursively for $k = 3, ..., L-3$ we will get that in order not to violate the contradiction assumption, the maximum is attained when $D_0 = D_1 = ... = D_{L-2}$. However, this is definitely wrong since from F_0 it is easy to see that $D_0 \leq 1/(1-\delta^2)$, and we can easily find a value that satisfies the constraints and is greater than $(L-1)/(1-\delta^2)$ (for example, the one derived from Equation (2) for $\delta^2 > 0$).

This proves that for the maximal $A_{L-1} = \sum_{i=0}^{L-2} D_i$ we must have $F_{L-2} = 0$ and $F_{L-3} = 0$. We can continue inductively by assuming $F_k = 0$ for $k = L-3, \ldots, 1$ and proving in the same way that $F_{k-1} = 0$.

We prove now that (ii)⇒(i). Assuming for all ℓ $F_l(D_0, ..., D_l) = 0$, we get the recursive formula from Equation (2). This formula defines $L-1$ linear equations that can be represented as a triangular matrix, i.e. they are linearly independent. This means there exists a solution and it is unique.

Now we show that (iii) $\Longleftrightarrow$ (i). Denote by $M = A_{L-1}$ the maximal value of $\sum_{i=0}^{L-2} D_i$ for a given δ. Equation (2) defines a function $M(\delta) = \sum_{i=0}^{L-2} D_i(\delta)$. It is easy to see from Equation (2) that $M(\delta)$ is a strictly increasing function of δ (since the D_ls increase as δ increases). Therefore, there exists an inverse function $\delta(M)$, which for any given value A_{L-1}, returns δ such that $M(\delta) = A_{L-1}$. This is the minimal δ since a smaller δ will not satisfy the relative error constraints, whereas a larger δ will not be minimal. ∎

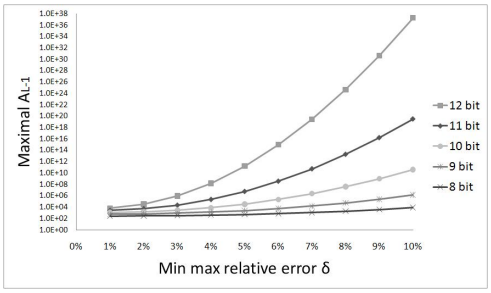

Fig. 4. The minimal max-relative-error, δ, vs. maximal estimator value, A_{L-1}, for several bit widths, $log(L)$

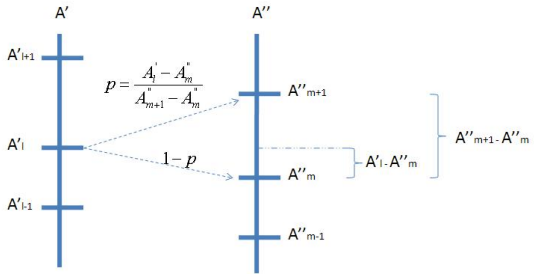

Fig. 5. CEDAR scaling - A' is a CEDAR estimation array that matches a relative error δ_0. A'' is a CEDAR array that matches a relative error $\delta_0 + \delta_{\text{step}}$, therefore it has an higher maximal value. Each flow pointer that points to an estimator in A' is converted to one of the two closest estimators in A'', according to probability p.

B. Capacity Region of Static CEDAR

Figure 4 illustrates the maximal possible guaranteed region for CEDAR, as provided by the min-max optimization. It plots the maximal estimator value as a function of the min-max relative-error, given several possible bit widths, as indicated in Equation (2). For instance, given $\log_2 L = 10$ bits, i.e. $L = 2^{10} = 1024$ estimators, we can achieve a min-max relative error of $\delta = 5\%$ using a maximal estimator value of $A_{L-1} = 32 \cdot 10^3$, and an error of $\delta = 10\%$ using $A_{L-1} = 35 \cdot 10^9$.

Alternatively, each plot also defines the lowest possible relative error guarantee given a counter scale, as proved in Theorem 4. For instance, given $A_{L-1} = 32 \cdot 10^3$, it is not possible to provide a better guarantee than $\delta = 5\%$.

Figure 4 also shows how an increase in the number of bits, i.e. in the available memory resources, leads to an increase in the maximal counter value and/or a decrease in the guaranteed relative error.

IV. Dynamic up-scale

A. Motivation

We now introduce the full CEDAR algorithm that can *dynamically adjust* to the counter scale. This algorithm is used to support an unlimited counter value, yet also provide a good relative error at any given time. The strength of our algorithm is the fact that we control the up-scaling by adjusting the relative error rather than changing the maximal estimation value. In this way the error of the entire counter scale is always under control and can be fine tuned by the user.

In practice, we start with the best estimation array that matches an initial maximal relative error δ_0, i.e. the estimation array with the largest maximal estimator. This estimation array is unique, as we proved in Theorem 4. As more and more packets are streaming in and the maximal flow counter is going to exceed our maximal estimation value, we set a new relative error of $\delta_{\text{new}} = \delta_0 + \delta_{\text{step}}$ and build a new matching estimation array. We further proceed in the same way for the next iterations.

B. Algorithm

In order to support up-scaling we use two estimation arrays A' and A'' in a simple *ping-pong* way. We initialize A' with equal-error estimators that match an initial error of δ_0 according to Equation 2, and likewise initialize A'' with estimators that match $\delta_0 + \delta_{\text{step}}$. Therefore A'' has an higher maximal estimation value than A'.

We start by using the static CEDAR algorithm with the A' array. Once the maximal flow pointer is getting close to the end of A' array such that $max_j(F_j) \geq i_{threshold}$, where $i_{threshold}$ is an arbitrary threshold value, we perform the following up-scale procedure: For j=0 to $N-1$ do

1) Set $l = F_j$
2) Find m such that $A''_m \leq A'_l < A''_{m+1}$ (e.g. using binary search within A'')
3) Set $p = \frac{A'_l - A''_m}{A''_{m+1} - A''_m}$
4) Set $F_j = m+1$ with probability p and $F_j = m$ with probability $1-p$

Figure 5 depicts a single step of the up-scale procedure.

We run the up-scaling procedure in parallel with the real-time counter updates, i.e. *without stopping* the CEDAR update algorithm. We use the current-flow index f as follows: assuming a new packet arrival for flow F_i, if $i < f$ we use the new estimation array A'', and otherwise use A' (since current flow f has not been up-scaled yet).

When the up-scaling procedure ends, we prepare A' for the next up-scaling event by filling it with new estimation values that match a new error of additional δ_{step}, i.e. after a new up-scaling event the relative error is $\delta_0 + 2 \cdot \delta_{\text{step}}$ and after n events it is $\delta_0 + n \cdot \delta_{\text{step}}$.

C. CEDAR Up-Scale Unbiasedness

In the following theorem we show the unbiasedness of CEDAR up-scale algorithm. Since we know from the unbiasedness of static CEDAR that estimation updates are unbiased as long as the scale is maintained, the proof essentially shows that up-scaling does not introduce bias. As for static CEDAR, the result assumes that either the estimator scale is unbounded, or up-scaling is done such that the maximum estimator value is never reached.

Theorem 5: The CEDAR up-scaled estimators are unbiased.

Proof: In the up-scaling phase, let C denote the real number of packets for the flow, $\hat{C}'$ its estimation under the old scale, and $\hat{C}''$ its estimation under the new scale. Then we need to show that $\mathbb{E}\left(\hat{C}''\right) = \mathbb{E}\left(\hat{C}'\right) = C$. Using the law

of iterated expectation $\mathbb{E}\left[\hat{C}''\right] = \mathbb{E}\left(\hat{C}''|\hat{C}'\right)$, it is therefore sufficient to show that $\mathbb{E}\left(\hat{C}''|\hat{C}'\right) = \hat{C}'$.

Consider now a random path of the counter estimator. As illustrated in Figure 5, denote by A'_l the random variable that represents the estimation value before the up-scaling. Also define A''_m on the new scale such that $A'_l \in [A''_m, A''_{m+1})$. Then we just need to show that for any such A'_l, $\mathbb{E}\left(\hat{C}''|\hat{C}' = A'_l\right) = A'_l$. Denote a Bernoulli random variable with parameter p as $B(p)$. Then we obtain:

$$\begin{aligned}\mathbb{E}\left(\hat{C}''|\hat{C}' = A'_l\right) &= \mathbb{E}\Big(A''_m + (A''_{m+1} - A''_m) \\ &\qquad \cdot B\left(\frac{A'_l - A''_m}{A''_{m+1} - A''_m}\right)\Big) \\ &= A''_m + (A''_{m+1} - A''_m) \cdot \frac{A'_l - A''_m}{A''_{m+1} - A''_m} \\ &= A'_l\end{aligned}$$

■

V. CEDAR Evaluation

We evaluate the performance of CEDAR on real Internet packet traces ([26] and [27]). First, we verify the unbiasedness of the CEDAR estimator, and also check that its relative error is constant over the entire counter scale, as proved in Theorem 4.

In Figure 6 we show the CEDAR estimation results with a 4KB estimation array (12-bit estimators). The initial scale is set to a relative error of $\delta_0 = 1\%$ with steps of $\delta_{\text{step}} = 0.5\%$. The maximal value we reached is 10^6 which is equivalent to an error of 3%. As it shown in the figure, CEDAR keeps its unbiasedness over the entire scale as well as its relative error.

In Figure 7 we compare the CEDAR relative error to SAC and DISCO with two different Internet traces. In Figures 7(a) and 7(b) we are using 12-bit estimators. CEDAR parameters are just the same as the previous experiment. For SAC we used a 8:4 bit allocation as recommended in [20], that is, 8 bits to the magnitude (which is denoted in [20] by $A[i]$) and 4 bits to the exponent (denoted by $mode[i]$). As DISCO introduces no up-scaling scheme we need to make some assumption on the maximal estimated value in order to compare it to SAC and CEDAR, we assume it can scale up to 32-bit counter. Therefore, we set the DISCO exponent parameter to $b = 1.0041$ which can reach the maximum real counter limit with no need to do up-scaling.

As shown in the figure SAC achieves the best error in small counter values, since for small values it can represent the exact counter value. However, from counter value of 10^3 and on we see that it performs the worst, reaching a relative error of above 5%. Regarding DISCO, it performs better than SAC, however it is bounded between 4% and 5%. As we can see, CEDAR keeps a near-constant relative error throughout the whole counter scale in spite of the fact that it performed 4 scaling steps. CEDAR relative error is better by third from that of DISCO and SAC. In Figure 7(c) we use trace [27] again but with 8-bit estimators. For CEDAR we use again $\delta_0 = 1\%$ and $\delta_{step} = 0.5\%$, for SAC we use a 5:3 bit allocation and for DISCO we use an exponent parameter of $b = 1.079$. As we can see again, CEDAR max relative error is better by third from both SAC and DISCO. Interestingly, we can see how SAC estimator restriction to the form $a2^b$ causes high oscillations in its relative error.

In order to check the error adaptation of the CEDAR up-scaling algorithm, we ran again trace [26] using 12-bit estimators but this time we stop the stream at three different points: after 1% of the trace, 10% and 100%. At every point we sample the counter estimations through the entire counter scale and measure the relative error. We compare the results to a static CEDAR, i.e. a CEDAR array that uses no up-scaling and relies on the maximal counter value from the beginning. Figure 8 shows the relative error results. We can see in Figure 8(a) that without the scaling algorithm the error is always 3% no matter what the real maximal counter value is. When using scaling, like in Figures 8(b) and 8(c), we can see that the error is adjusted to the current maximal counter (after 1% and 10% of the trace).

VI. FPGA Implementation

A. Implementation Description

In this section we examine CEDAR implementation using Xilinx Virtex-5 FPGA. We use 192KB RAM for the flow pointer array and 16KB for the estimation array. Each estimator is 32 bit width. The estimation values (given by Equation(2)) are multiplied by 1000 in order to avoid float operations, yet keep a reasonable accuracy. The maximal design clock rate is 170MHz, therefore, assuming packets are accessed through a 32-bit bus with the same clock, our design may support streams of up to 5.4Gbps. Without any optimization our design uses 2500 1x6 LUT and 2500 flip-flops (roughly equivalent to 12K gates on ASIC design). In order to enable fast reading and writing operations we use DPR (Dual Port RAM) for all arrays, in this way both the CEDAR module and the client application can access any array at any time. The CEDAR module has an I/O register interface in order to program it with initial configurations, e.g. array sizes, and commands, e.g. start/stop.

Figure 9 depicts our implementation block scheme. The CEDAR module itself is built of a simple state machine. For any incoming packet we first fetch the flow pointer, then we fetch the two successive estimators and finally we decide whether to update or not the flow pointer. Both the flow pointer array and the estimation array are stored in a dual-port RAM in order to enable the client application fast reading of the flow counter estimation at any given time.

A second estimation RAM is drawn with a dashed line in order to support the up-scaling algorithm. The up-scaling algorithm is done by a dedicated SW application. In order to trigger this application, CEDAR generates an interrupt when the maximal flow index exceeds a programmable threshold. In addition, CEDAR has a *current-up-scaled-flow-index* register. This register is updated by the up-scaling application during its iteration through the flow pointer array. CEDAR uses this register in two manners: to decide which estimation array should be used for a specific flow (by comparing the flow pointer to that register) and to lock the updating of the current up-scaled flow.

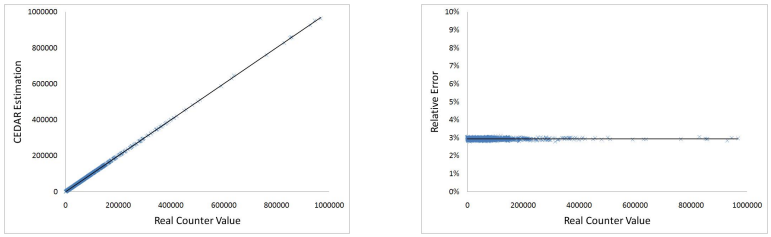

(a) CEDAR unbiased estimation vs. real counter values (Property 1)

(b) CEDAR constant relative error (Theorems 3 and 4)

Fig. 6. CEDAR results on a a real IP packet trace [26] using 12-bit estimators

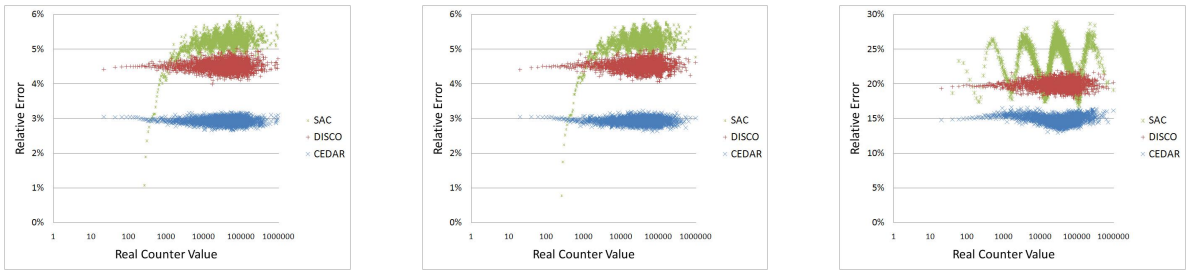

(a) Trace [26] counter estimation using 12-bit estimators

(b) Trace [27] counter estimation using 12-bit estimators

(c) Trace [27] counter estimation using 8-bit estimators

Fig. 7. Relative error comparison on real IP packet traces [26] and [27] using 12-bit and 8-bit estimators

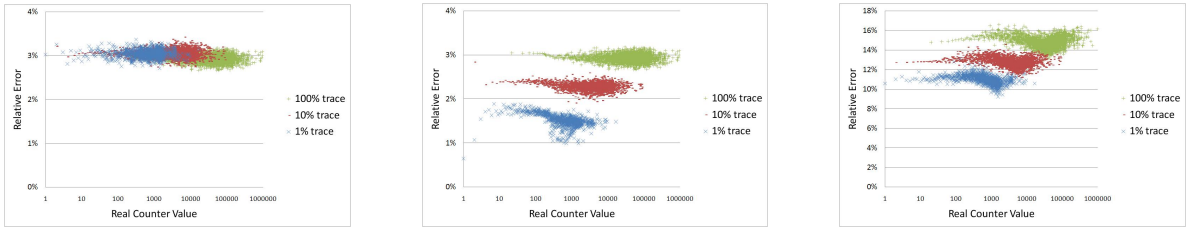

(a) CEDAR without scaling for 12-bit estimators. The maximal value is aimed at 10^6. The matching relative error is kept at 3% throughout the trace

(b) CEDAR with scaling for 12-bit estimators. The relative error is adjusted according to the current maximal counter value

(c) CEDAR with scaling for 8-bit estimators. The relative error is adjusted according to the current maximal counter value

Fig. 8. Relative error comparison between non-scaling CEDAR to a scaling CEDAR using trace [26]

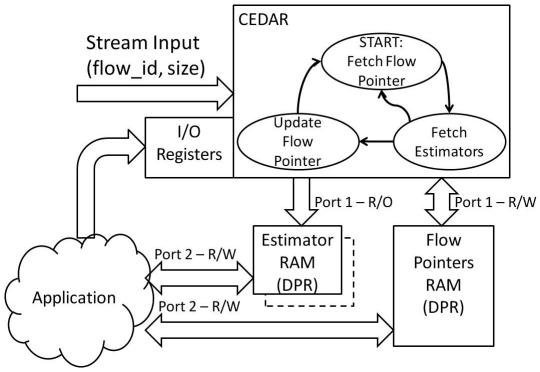

Fig. 9. CEDAR FPGA implementation block scheme - CEDAR is implemented as three states SM. We use DPR (Dual Port RAM) for estimator array and for flow pointers. Thus, both CEDAR module and application can access arrays to perform fast updating and fast reading

B. Performance Analysis

We would like now to analyze the proposed model performance. As shown in the figure, we have 3 RAM accesses and sometimes even 4: one reading of the flow pointer, two readings of successive estimators and an optional writing of a new flow pointer. In our implementation we handle each packet separately. That is, we do not accept the next packet till we finish a full cycle of the CEDAR state machine. This imposes a delay between 2 successive incoming packets of at least 4 cycles. Typically, this is not a real problem since the packet minimal size is bounded by its header size which is usually longer than 4 words. So assuming we have a word access per cycle, our 4-cycle delay will not cause any stream halt.

To conclude, we get a simple implementation which can work with high stream bit-rates and enable fast reading by the client application. Assuming an ASIC solution is much faster than FPGA, our implementation can probably support up to tens of Gbps. In terms of size, our implementation requires in the worst case 12K Gates, which is negligible comparing to the flow pointer RAM requirement.

VII. Conclusion

In this paper we proposed a simple counter estimation method called CEDAR. We proved that it can guarantee optimal estimation in the manner of *min-max-relative-error*. We proved that the minimal boundary is attained when all estimators have *equal-relative-error*. In addition, we introduced a simple up-scaling algorithm which enables us to adjust the relative error of the entire counter scale according to the current maximal counter without having any pre-knowledge about the flow distribution. We also proved the unbiasedness of the CEDAR up-scaling algorithm.

Then we demonstrated the CEDAR estimation using two different Internet traces. We showed its unbiasedness as well as relative error equality through the entire counter scale. We also showed that it achieves the min-max boundary, up-scales smoothly, and adjusts the relative error estimation to the maximal flow counter.

Last, we showed a simple CEDAR implementation on FPGA that enables the client application to perform a fast reading of counter estimation at any given time. Our implementation consumes a relatively small space for the CEDAR logic part and can support high bit rates. We synthesized our FPGA to support a bit stream of 5.4Gbps, and expect an equivalent ASIC implementation to support tens of Gbps.

Acknowledgement

The authors would like to thank William Backshi for the CEDAR implementation on FPGA. The work was partly supported by the European Research Council Starting Grant n°210389.

References

[1] Cisco, CRS-1 multishelf system. [Online]. Available: http://www.cisco.com/en/US/products/ps5842/

[2] Juniper, TX matrix plus. [Online]. Available: http://www.juniper.net/us/en/products-services/routing/t-tx-series/txmatrix-plus/

[3] Broadcom-Dune, FE600 fabric element. [Online]. Available: http://www.dunenetworks.com/webSite/Modules/News/NewsItem.aspx?pid=355n&id=89g

[4] F. Abel, C. Minkenberg, I. Iliadis, A. P. J. Engbersen, M. Gusat, F. Gramsamer, and R. P. Luijten, "Design issues in next-generation merchant switch fabrics," *IEEE/ACM Transactions on Networking*, vol. 15, no. 6, pp. 1603-1615, 2007.

[5] A. Kabbani, M. Alizadeh, M. Yasuda, R. Pan and B. Prabhakar, "AF-QCN: Approximate fairness with quantized congestion notification for multi-tenanted data centers," *18th IEEE Hot Interconnects*, Mountain View, CA, August 2010.

[6] G. S. Manku and R. Motwani, "Approximate frequency counts over data streams," *28th International Conference on Very Large Data Bases*, 2002.

[7] X. Dimitropoulos, P. Hurley and A. Kind, "Probabilistic Lossy Counting: An efficient algorithm for finding heavy hitters," *ACM SIGCOMM Computer Communication Review*,vol. 38, no. 1, January 2008.

[8] L. Fan, P. Cao, J. Almeida, and A. Z. Broder. Summary cache: a scalable wide-area web cache sharing protocol, *IEEE/ACM Transactions on Networking*, 8:3, pp. 281–293, 2009.

[9] C. Estan, G. Varghese. New directions in traffic measurement and accounting," *ACM Transactions on Computer Systems*, 21(3), 2003.

[10] Y. Lu, A. Montanari, B. Prabhakar, S. Dharmapurikar, A. Kabbani, "Counter braids: a novel counter architecture for per-flow measurement". *SIGMETRICS*, pp. 121–132, 2008.

[11] Y. Lu, B. Prabhakar, "Robust counting via counter braids: an error-resilient network measurement architecture," *IEEE Infocom*, pp. 522–530, 2009.

[12] N. Hua, B. Lin, J. J. Xu, and H. C. Zhao, "Brick: A novel exact active statistics counter architecture," *ANCS*, 2008.

[13] D. Shah, S. Iyer, B. Prabhakar, and N. McKeown, "Maintaining statistics counters in router line cards," *IEEE Micro* 22(1): 76-81 2002.

[14] S. Iyer, R. R. Kompella, and N. McKeown, "Designing packet buffers for router linecards," *IEEE/ACM Trans. on Netw.*, 16(3): 705-717, 2008.

[15] M. Roeder and B. Lin. "Maintaining exact statistics counters with a multilevel counter memory," *IEEE Globecom*, Dalas, USA, 2004.

[16] Q. Zhao, J. Xu, and Z. Liu. "Design of a Novel Statistics Counter Architecture with Optimal Space and Time Efficiency," *ACM SIGMETRICS*, France, 2006.

[17] B. Lin and J. Xu, "DRAM is plenty fast for wirespeed statistics counting," *ACM Performance Evaluation Review*, vol. 36, no. 2, pp. 45-50, September 2008.

[18] B. Lin, J. Xu, N. Hua, H. Wang, H. Zhao, "A randomized interleaved dram architecture for the maintenance of exact statistics counters," *ACM SIGMETRICS*, Seattle, WA, June 2009.

[19] H. Zhao, H. Wang, B. Lin, and J. Xu, "Design and performance analysis of a dram-based statistics counter array architecture," *ANCS*, Princeton, NJ, October 2009.

[20] R. Stanojevic, "Small active counters", *IEEE Infocom*, 2007.

[21] C. Hu, B Liu, and K Chen, "Discount counting," *IEEE ICNP*, 2009.

[22] P. Flajolet, "Approximate counting: a detailed analysis," *BIT*, vol. 25, pp. 113-134, 1985.

[23] E. W. Weisstein, "Quadratic curve," MathWorld. http://mathworld.wolfram.com/QuadraticCurve.html

[24] E. Cohen, R. F. Riesenfeld and G. Elber, "Geometric modeling with splines: an introduction", 2001.

[25] H. W. Kuhn, and A. W. Tucker, Nonlinear Programming, 1951.

[26] CAIDA Anonymized 2008 Internet Trace equinix-chicago 2008-03-19 19:00-20:00 UTC Direction A (DITL) (collection), http://www.caida.org/data/monitors/passive-equinix-chicago.xml, 1 hour trace on a single direction of an OC192 backbone link

[27] CAIDA Anonymized 2008 Internet Trace equinix-chicago 2008-03-19 19:00-20:00 UTC Direction B (DITL) (collection), http://www.caida.org/data/monitors/passive-equinix-chicago.xml, 1 hour trace on a single direction of an OC192 backbone link