

Bloom Filter for Double Counting Avoidance in Radio Frequency Ray Tracing

Roman Novak, *Member, IEEE*

Abstract—Brute force radio frequency ray tracing scales well on multicore computing architectures as each ray can be traced independently. The discrete nature of rays with no thickness shows a weakness in the aggregation step where nearby rays to the observation point need to be detected and differentiated. The fact that angular defect cannot be distributed evenly for more than twelve rays in space leads to a double counting phenomenon, i.e., the radius of a reception sphere cannot exclude all but one ray per wavefront. This either leads to significant signal errors or requires the use of space and time consuming wavefront differentiation. Bloom filters configured with marginal false positive rate are proposed here as a replacement of the exact wavefront differentiation, effectively eliminating double counting errors. Substantial space gains while computing channel impulse responses are reported. Further, due to the importance of ray-launching template grids for the double counting avoidance, the analytical angular bounds for the two common icosahedral grids are presented. The frequently referenced approximate double counting probability of 20.9% is shown to be highly underestimated when affected by the refraction- and diffraction-induced changes in ray spacing.

Index Terms—Algorithm optimization, Bloom filter, channel impulse response, radiowave propagation, ray tracing.

I. INTRODUCTION

RADIO frequency (RF) ray tracing is a valuable tool in the development of wireless technologies and services. It allows far superior modelling of electromagnetic propagation effects than any other stochastic modelling. Accuracy-wise it can only be surpassed by approximations of spatial and temporal derivatives appearing in Maxwell's equations, which is computationally infeasible if the wavelength is exceedingly small compared to the size of modelled environment. Advanced ray-tracing techniques take into account the majority of paths the real wavefront would traverse and model actual physical phenomena responsible for propagation of electromagnetic waves.

Many ray-tracing algorithms have been proposed since the first ideas by Ikegami et al. [1] in 1991. From the ray handling perspective, most of them fall into three computationally distinct groups. The first, often seen as the brute force approach, effectively traces a large number of rays from the transmitting source in all directions into the scene. The concept of a reception sphere is usually needed to detect rays passing by the receivers [2]. The algorithms from this group refer to the principle as ray launching [3], pincushion method [4] or ray shooting and bouncing (SBR) [5], which is also a designation used here. The second approach aggregates traced rays as

ray tubes [6] or beams [7] in order to reduce computational complexity. Finally, the third imaging approach [8] goes the furthest by using entire scene surfaces as ray aggregation units.

Here we restrict ourselves to the problem of double counting avoidance that emerges in the SBR group of algorithms [2]. Namely, signal aggregation at given observation point dictates the use of variable-sized reception spheres to detect at least one ray per passing wavefront. The geometry of ray distribution in space does not allow perfect detection of exactly one ray per wavefront. Double or even multiple hits need to be filtered out before signal accumulation.

Several approximate algorithms have been proposed, including [2], [3], [9], [10], as well as the exact double counting avoidance [11] at high space and time cost. Imperfect double counting may be tolerated for some simulation tasks, such as signal coverage estimation or even more demanding RSSI evaluation. Other tasks strive to eliminate algorithmic errors as much as possible. Most notable among them is that of finding accurate channel impulse responses (CIRs), with the need for precise simulation tools in the fields such as ultra-wideband communications and indoor localization.

The probability of double hits and the reception sphere sizes depend on the properties of icosahedral grids used as ray generation templates. We first review new findings on angular bounds that affect the double counting problem [12]. Next, predicated on the assumption of rays hitting the reception sphere at no particular order, the optimal wavefront differentiation is studied. Evaluating responses at multiple points simultaneously in a single simulation run while using exact wavefront differentiation turns out to be space prohibitive for many computing architectures. A straightforward reduction of excessive space requirements is to introduce some form of data compression without further prolonging the already time-consuming ray-tracing procedures. Apparently conflicting requirements can be satisfied if we allow some errors and keep the number sufficiently low. Bloom filters provide such space-efficient probabilistic data structures while offering constant-time update and set membership tests [13]. Their use as near-optimal wavefront differentiator is proposed here. Trade-off between size and accuracy is configured by false positive rate, which is filter's key parameter. Note that many applications in computing and particularly in communication systems make use of Bloom filters. Load balancing, routing, web caching, network security [14], databases and distributed storage systems [15] are just some of them, where managed data sets consist of firewall addresses, routing prefixes, TCP flow labels, cache indices, etc. Apart from the basic filter, numerous variants adapted to specific problem domains have been developed, with several of them introduced throughout

This work was supported by the Slovenian Research Agency under Grant No. L2-7664 and P2-0016.

R. Novak is with the Department of Communication Systems, Jožef Stefan Institute, SI-1000 Ljubljana, Slovenia (e-mail: roman.novak@ijs.si).

the paper. A comprehensive survey on the subject is available in [16].

Although the source of the double counting errors is clearly within the ray-tracing algorithm, we evaluate proposed algorithm also in the context of path loss measurements. Note, however, that a number of factors influence such a comparison, including, but not limited to, how well we capture the geometry, composition of materials, or presence of vehicles and people.

In summary, the contributions of this paper are as follows:

- Near-optimal double counting avoidance at significantly reduced space cost and lower time complexity than any known solution is proposed by the introduction of versatile Bloom filtering as a wavefront differentiation technique.
- Analytical bounds on the reception sphere radius are provided for the two common icosahedral grids used as the ray launching templates. Bounds suggested in the literature are vague, usually quoting approximate average separation angle.
- Irregularities in wavefront ray spacing are observed for refracted and diffracted rays if rays are traced in full compliance with Snell's law and geometrical theory of diffraction is employed.
- The effect of double counting is studied in the context of synthetic channel impulse responses and evaluated against measured path losses. Near-optimal double counting avoidance is shown to be mandatory when major sources of algorithmic errors need to be eliminated. Indoor and outdoor cases are presented.

In the following, Section II studies the reception sphere radius requirements for the two grid launching templates in the light of recent analytical findings. Refracted and diffracted rays are shown to further worsen the problem. Related work on the double counting avoidance is presented in Section III. Wavefront differentiation using Bloom filter is described in Section IV. The evaluation of the proposed filtering in the contexts of plane earth model, narrowband channel impulse responses and measured path losses is given in Section V. The use of more elaborate filters is discussed in Section VI, followed by the conclusion in Section VII.

II. DOUBLE COUNTING PROBLEM

The shooting and bouncing rays (SBR) effectively traces a large number of rays from the transmitting source in all directions into the scene while modelling surface reflections, through-wall transmissions, edge diffractions and possibly scattering phenomena. Using geometrical optics concepts in describing radio frequency propagation implies that initial rays are an abstraction of a single wavefront spreading into space. Subsequent electromagnetic interactions with matter initiate new wavefronts, described by another sets of reflected, refracted, diffracted or scattered rays. The signal evaluation at given observation point combines these wavefronts freely, in the same way as if they are being transmitted by multiple independent sources. However, each wavefront should be considered once and only once while passing by the observation point.

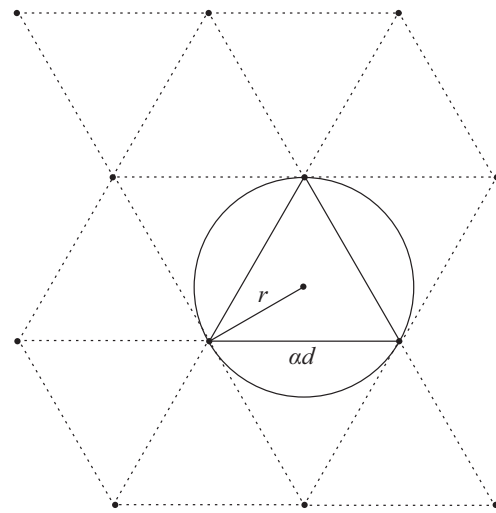


Fig. 1. Uniformly distributed rays on a wavefront surface are shown as vertices of equilateral plane triangles. In order to guarantee at least one ray hit, the minimum reception sphere radius is that of a triangle-circumscribed circle.

Conceptually, rays have no thickness and common approach to detect a nearby ray and thus a wavefront is to introduce a sphere object with non-zero radius and inspect the intersecting rays. Here the double counting problem emerges. Even under the unrealistic assumption of uniformly distributed initial rays it is not possible to have a reception sphere size that would guarantee one hit per wavefront while completely avoiding double hits. In Fig. 1 a wavefront surface is shown as a plane of dotted equilateral triangle grid with each triangle vertex representing a perpendicular ray. Adjacent vertices are uniformly spaced by αd , where α denotes a constant angular distance or separation between rays in radians and d the wavefront propagation distance. In order to enclose at least one vertex, a randomly drawn circle over the triangle grid should have a minimum radius of the triangle-circumscribed circle, i.e.,

$$r = \frac{1}{\sqrt{3}} \alpha d. \quad (1)$$

The wavefront surface is generally not a plane. Even though, the minimum reception sphere radius converges to (1) as α gets smaller. In this ideal scenario of uniformly distributed vertices, the probability of randomly positioned circle spanning more than one vertex is $2/\sqrt{3} (\pi/3 - \sqrt{3}/2)$ or approximately 20.9% [2], which is frequently cited double counting probability.

The above value is only a lower limit for a number of reasons. First, a uniform distribution of rays in three dimensions cannot be achieved for more than 12 rays and maximum α should be used in (1). The maximum could be local over adjacent rays, which is difficult to implement, or global over all pairs of adjacent rays. Secondly, the refracted and the diffracted rays change spacing between each other in a non-trivial way, making theoretical double counting probability harder or even impossible to achieve.

Icosahedral grids have been used for launching the initial set of rays since early SBR proposals. The maximum angular separation of grid vertices has been largely left to the numerical




Fig. 2. Recursive versus non-recursive icosahedral grid construction; selected base triangle (a) is recursively divided and projected twice (b, c), whereas non-recursive construction starts with the final number of equilateral partitions (d).

evaluation. Only recently the research of geometric properties of the icosahedral-hexagonal grids offered analytical bounds for several grid construction algorithms [12]. The reception sphere radii in the presented simulations are established using those bounds.

A. Rays Maximum Angular Separation

Icosahedral grid on a sphere around the transmission point provides excellent template for the initial set of rays. According to Gauss-Bonnet theorem, icosahedron distributes the angular defect evenly over the most number of vertices among all regular polyhedrons. Consequently, 12 rays defined by 12 icosahedron vertices on the surface of the enclosing sphere is the maximum number of rays in space with constant angular separation. Icosahedral grids, which are obtained by dividing the initial 20 congruent equilateral plane triangles and projecting denser triangle mesh to the enclosing sphere surface, are commonly used to generate larger number of rays with rather uniform distribution.

The recursive icosahedral grid construction is the most straightforward and, due to its simplicity, commonly used by the researchers. Each recursive step halves the sides of plane triangles and projects newly introduced points on the enclosing sphere to quadruple the number of spherical and corresponding plane triangles. The refinement level n establishes the final $10 \times 2^{2n} + 2$ grid points acting as initial rays directions. In Fig. 2, selected base triangle (a) is first divided into four triangles with the vertices on the enclosing sphere (b). The following recursive step doubles the number of plane triangles (c).

On the other hand, the non-recursive construction has greater flexibility in controlling the number of grid points. The algorithm subdivides the initial icosahedron plane triangles

into a number of smaller congruent equilateral triangles, as illustrated in Fig. 2 (d), and projects them on the enclosing sphere in a single step. If initial triangle side is partitioned in m sections, the construction produces $10 \times m^2 + 2$ grid points.

Rays angular distribution differs with respect to the chosen grid construction, thus affecting the double counting differently. Whereas the initial plane triangles are equilateral, this is not the case at higher refinement levels. For a recursive construction, the exact relation between rays maximum angular distances at the subsequent refinement levels is obtained from the proof of Lemma 3.12 given in [12] as

$$\sin \frac{\alpha_{n+1}}{2} = 2^{-n} \sin \frac{\alpha_n}{2} / \cos \frac{\alpha_n}{2}. \quad (2)$$

Using the initial angular separation of icosahedral grid points measured from the icosahedron center

$$\alpha_0 = 2 \sin^{-1} \sqrt{\frac{2}{5 + \sqrt{5}}} \approx 63.4349 \text{ deg} \quad (3)$$

one can calculate the exact maximum angular separation α_n at refinement level n .

For the non-recursive grid construction only the upper bound on the maximum separation angle is analytically known, however the bound is approached fast. Given the initial subdivision of m sections, the maximum angular distance α_m converges to

$$\begin{aligned} \lim_{m \rightarrow \infty} \alpha_m &= \frac{1}{m} \sin\left(\frac{2\pi}{5}\right)^{-1} \left(\cos^2 \frac{\alpha_0}{2} - \tan^2 \frac{\pi}{6} \sin^2 \frac{\alpha_0}{2}\right)^{-\frac{1}{2}} \\ &\approx \frac{75.812}{m} \text{ deg}, \end{aligned} \quad (4)$$

as shown in the proof of Theorem 3.14, also provided in [12].

For the purpose of interpreting paper results it is worthwhile to mention that the ratio between the longest and the shortest angular distance increases with the grid refinement level, and it converges to approximately 1.195114 and 1.479348 for the recursive and the non-recursive grids, respectively [12]. In general, the non-recursive construction produces larger angle variations.

B. Ray Spacing

Equation (1) is a tight bound for the line-of-sight and reflected rays that guarantees at least one hit per wavefront, given maximum separation angle α and propagation distance d . Such a property does not hold for the refracted as well as the diffracted rays in general.

If refracted rays are traced through general-shaped obstacles in compliance with Snell's law, spacing between rays is affected by the obstacle material properties and angles of incidence. Both increased and decreased spacing may occur after rays leave the obstacle. However, when the environment is composed of only wall-shaped objects with parallel interface planes, which is commonly the case in indoor architectures, equation (1) remains an upper bound for the reception sphere size. Rays entering a transmission medium with higher index of refraction, i.e., $n_2 > n_1$, change directions and thus angles between adjacent rays in compliance with $\theta_{tr} =$

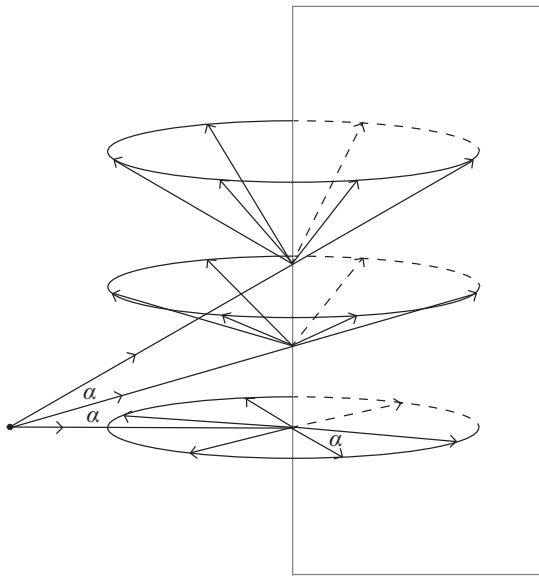


Fig. 3. Regular spacing of initially launched rays is distorted if edge diffraction is modelled using Keller cones. Rays on a cone surface are spaced closer than rays on adjacent cones when hitting a reception sphere.

$(n_1/n_2) \sin \theta_{inc}$, where θ_{inc} is the ray's angle of incidence and θ_{tr} is the angle of transmission. On leaving the transmission medium with higher index of refraction at the opposite parallel interface and reentering the initial transmission medium the angle between the two rays is restored while leaving rays closer to each other than expected by (1). In principle, equation (1) can be rewritten by taking into account the indices of refraction for both mediums and the obstacle depth. As the calculation is non-trivial with possibly significant impact on the ray-tracing performance, additional hits are commonly allowed. Therefore, in a simulation, double, triple or even more hits need to be filtered out by a double counting avoidance algorithm.

The double counting problem becomes even worse if edge diffraction phenomenon needs to be simulated. Diffraction can be avoided in many indoor scenarios but not in the outdoor urban scenarios where major signal contributions originate from the diffracted paths. Modelling edge diffraction in the SBR group of algorithms is commonly done using the geometrical theory of diffraction [17] as illustrated in Fig. 3. Wavefront rays passing in close vicinity to a diffraction edge produce a number of diffracted rays on the surfaces of Keller cones. Maximum separation angle α is typically used to distribute rays around the cone surface because adjacent cones are already α -spaced. The reception sphere size (1) should match rays from different cones, thus the distance to the transmission source is used. However, rays on the same cone surface originate closer to the reception sphere and introduce additional hits than predicted by (1).

III. DOUBLE COUNTING AVOIDANCE

Erroneously updating signal due to a double hit leads to +6dB power error per ray. One of the first attempts to reduce the number of double hits is so-called Zero Counting (ZC) algorithm [9]. Instead of (1) half the distance between rays is

proposed for the reception sphere radius, which results in a number of missed wavefronts. Some of them are captured by introducing auxiliary rays at the center of the icosahedral grid faces, for which smaller $\sqrt{3}/3 - 1/2$ times the rays distance is suggested for the reception sphere radius. Zero count probability is estimated to be 4.97% in ideal case of uniformly distributed rays. However, reduced spacing between rays due to the strict interaction modelling is entirely overlooked. The SBR that employs ZC could still experience double hits while the total number of launched rays needs to be tripled.

A weighting approach to the double counting problem is proposed in the Distributed Wavefronts (DW) algorithm. All rays hitting the reception sphere contribute a limited amount of signal power. The contribution is weighted in a way that roughly produces signal levels invariant of the local distribution of rays [2]. Although the proposed weighting function provides a useful compromise with low processing overhead, it cannot replace the exact solution. The weights are based on a Monte Carlo simulation designed to obtain best average fit. Even though the algorithm is not a classical double counting avoidance, it serves as one of the reference algorithms.

The Characteristic Sequences (CS) algorithm provides exact double counting avoidance at high space and time cost [11]. Each wavefront is uniquely identified by a sequence of scene interactions that rays experience before hitting the reception sphere. Interactions with the same planar surface or with the same diffraction edge are considered identical. Such a sequence is called a characteristic sequence. In order to distinguish rays belonging to different wavefronts—and reject those representing already seen wavefronts—their characteristic sequences are constructed while advancing rays through the scene. A set of sequences is memorized per observation point. On each new hit, ray's sequence membership is tested. If test fails, the newly arrived ray acts as a valid signal carrier and its characteristic sequence is added to the set. Since the sets are generally expected to be small, at least for typical use cases, they are commonly stored as tables. A membership query is then of $O(n)$ time complexity, with constant update cost. Here we use n to designate the expected number of wavefronts at the observation point. Critical in most ray-tracing implementations is the algorithm storage requirement. Let *depth* be the maximum number of scene interactions a ray is allowed to encounter during a simulation. Then the set of characteristic sequences is $O(n \times \text{depth})$ space bound.

Proposals of more compact representation of lengthy characteristic sequences emerged early in the SBR research. Replacing a sequence with a pair of ray's incident angle and interaction count is rationalized in [10]. The excessive memory requirement reduces to linear dependence on n , whereas the set membership tests remain of the same time complexity but with smaller constant factor. The double counting avoidance algorithm, which we denote as the Angle and Interactions (AI) algorithm, proclaims two rays with the same number of interactions and similar sphere hitting directions as belonging to the same wavefront. This is clearly a heuristic solution. There is a possibility that two different wavefronts come from similar directions while experiencing equal number of scene interactions. Further, the angle similarity criterion is

vaguely defined, as it is further explained in the evaluation section. Both, time and space complexities of the AI algorithm are $\mathcal{O}(n)$ while permitting some double counting errors for smaller memory footprint.

Similar space and time observations hold for the Angle and Distance (AD) algorithm, in which the number of interactions is simply replaced with the propagation distance [3]. The algorithm is even more imprecise than AI. The fact that even rays of the same wavefront experience slightly different refraction angles and consequently path lengths makes differentiation based on the angles and distances highly susceptible to a membership error.

IV. BLOOM FILTERING

As presented so far, the only acceptable wavefront differentiation option is some variant of error-free characteristic sequences, the use of which is limited by high memory requirements. The aim of the proposed Bloom filtering is to significantly reduce the needed space and time while providing nearly optimal double counting avoidance. As shown in the next section, the improvements can be remarkable.

Bloom filter emerged several decades ago as a lossy-encoded hash table allowing some errors in membership testing while allocating less memory than needed for the lossless data structure [13]. The underlying idea of the proposed double counting avoidance algorithm is to configure Bloom filter with a small but still tolerable error rate while taking the advantage of its inherent dataset compression.

Formally, Bloom filter is a data structure representing a set of elements $S = \{x_1, x_2, x_3, \dots, x_n\}$ and supporting at least two operations—set addition and a membership query. Plain array of m bits, denoting the filter size, is used to store membership information instead of buckets or slots of a more conventional hash table. Initially all bits are set to 0.

In order to add element x to set S , k independent and preferably uniform hash functions $h_i(x), 1 \leq i \leq k$, are first calculated. Bits at matching indices are then set to 1. The membership query of a basic Bloom filter consists of testing corresponding k bits. If all bits are set, the element is assumed to be a member. It follows that all elements actually added to the set are correctly recognized. On the other hand, false positives cannot be ruled out, meaning that some elements are erroneously claimed to be in the set even though they have never been added.

The false positive probability p , also called false positive rate, is a function of the filter size m , of the number of hash functions k and of the number of elements n added to the set. The widely used equation linking these parameters to the false positive rate

$$p = (1 - (1 - 1/m)^{kn})^k \quad (5)$$

has been shown to be too optimistic due the incorrect probability expectation of a single bit to be one [18]. The functional dependence is more closely described as

$$p = \frac{m!}{m^{k(n+1)}} \sum_{i=1}^m \sum_{j=1}^i (-1)^{i-j} \frac{j^{kn} i^k}{(m-i)! j! (i-j)!} \quad (6)$$

However, significant disparities can be observed only for extremely small filter sizes of few bytes in length and large m/n ratios. Therefore, it is acceptable to approximate (5) by $(1 - e^{-kn/m})^k$ and analytically derive the optimal number of hash functions of larger filters as

$$k_{\text{opt}} = \frac{m}{n} \ln 2. \quad (7)$$

Given false probability rate, number of added elements and optimal number of hash functions the required filter size is then

$$m = -\frac{n \ln p}{(\ln 2)^2}. \quad (8)$$

A. Filter Parameters

Bloom filter as an index of wavefronts hitting a reception sphere reduces overall double counting avoidance memory requirements, with additional benefits of the constant build and query times. The double counting rejection efficiency of a filter should be close to the optimal wavefront differentiation. In order to avoid the majority of erroneous rejections of rays at an observation point, the false positive rate should be kept low. Based on the evaluation results provided in the next section, and for the purpose of the channel impulse response calculations, we propose choosing the probability p less or equal to 10^{-4} indoors and less or equal to 10^{-2} outdoors.

Fixed size Bloom filter in a pre-allocated memory space is best option for a fast ray-tracing implementation. Thus, a bound on the number of elements n should be known in advance. Because the number of different wavefronts at given observation point greatly varies with the simulation parameters, one could get a good estimate by a test run, in which only reception sphere hits are counted. The number can then be corrected by a double count probability, giving an upper bound on the number of elements, and by using (8), the filter size. Note that due to the decreasing ray spacing, as discussed earlier, the assessed number of elements would probably be higher than the actual number of different wavefronts.

Filter size estimation can be eliminated by the use of scalable Bloom filters [19], which gradually increase storage capacity while assuring given false positive rate. Conceptually different and architecturally more demanding dynamic Bloom filters [20] offer a solution as well if the set cardinality is unknown a priori. The third option is the use of split Bloom filters [21], which initially start with more space than projected and rebuild themselves if false positive rate is reached. However, all these approaches are not well-suited to the ray-tracing domain where a ray is commonly handled by a lightweight thread, possibly executing on massively-parallel hardware. Nevertheless, more elaborate solution is a topic of future research.

B. Hashing

Hash functions are at the core of Bloom filters. The previously stated relations between filter parameters are based on the assumption of mutually independent and uniformly distributed hash outputs. Numerous complex functions can

satisfy the criteria. However, when applied to the ray-tracing domain, the cost of computing a single hash should be as low as possible. This excludes not only highly complex cryptographically strong hashes, such as SHA1 or MD5, but also other less perfect but still time consuming functions. Favorably, it has been shown that randomness in hashed input greatly assists simple hash functions in achieving high level of uniformity [22].

For the purpose of this research, a linear congruential generator, which is commonly used as a pseudorandom number generator, is combined with characteristic sequences introduced earlier. Let each scene surface and each diffraction edge have a unique id . A wavefront characteristic sequence is then a concatenation

$$x = id_1 || id_2 || id_3 || \dots || id_{depth}. \quad (9)$$

Further, let

$$\text{lcg}(seed) = (a \times seed + c) \bmod b \quad (10)$$

be a 64-bit implementation of common linear congruential generator with $a = 6364136223846793005$, $c = 1$ and $b = 2^{64}$. The generator is a variant of Knuth's MMIX code that can be found in many Linux C libraries.

First, we iteratively chain outputs of pseudorandom number generator, seeded with interaction identifiers, to obtain base hash value

$$H_1(x) = \text{lcg}(\text{lcg}(\text{lcg}(id_1) \oplus id_2) \dots \oplus id_{depth}), \quad (11)$$

where \oplus stands for bitwise addition modulo 1. We define a sequence of the next $k - 1$ pseudorandom numbers as the remaining hash values

$$H_i(x) = \text{lcg}(H_{i-1}(x)), \quad 2 \leq i \leq k. \quad (12)$$

The indices into filter's bit array are then remainders of the division of the upper 32-bits of the generator state by filter size m

$$h_i(x) = (H_i(x) \gg 32) \bmod m, \quad 1 \leq i \leq k. \quad (13)$$

Note that using the upper bits of a pseudorandom generator state is a common practice to avoid short periods in the low-order bits when modulus b is power of 2.

The generator has 64-bit internal state with a period of $P \leq 2^{64}$. In order not to fail random goodness tests, no more than \sqrt{P} numbers should be used from a linear congruential generator [23]. Consequently, the number of different wavefronts at each observation point should not exceed $2^{32}/(depth + k - 1)$. The above hash formulation effectively supports Bloom filters of up to 4 Gbits in size at the cost of one addition and one multiplication per ray interaction.

V. CHANNEL IMPULSE RESPONSE EVALUATION

Rigorous double counting avoidance is important for advanced channel characterization, including but not limited to delay spread, direction of signal arrival and channel impulse response. Since channel impulse response (CIR) completely describes a linear time-invariant communication channel at

all frequencies and efficient CIR measurements are becoming intrinsic part of pulse based physical layers, the need for precise CIR simulation is on rise. For example, the research of UWB indoor localization with the CIR measurements at the core of the technology is expected to increasingly rely on the use of simulation tools. Another example is ray-tracing based fingerprint positioning technique in OFDM networks [24]. CIR simulations in urban scenarios have also been long recognized as an important tool in the design of transmission systems [25].

In the following, the performance of the proposed Bloom filtering is evaluated in the CIR context. First, the quality of evaluated channel impulse responses with respect to double counting avoidance algorithms is assessed. Space requirement, being the most limiting parameter of CIR simulations, is considered to be the second most important metric. Optimal CS provides exact solution against which other algorithms are evaluated.

Channel impulse response is typically modeled as a sum of time-varying number of multipath components in a tap-delayed linear filter. Time-invariant model is adopted in this research as time-varying aspects of channel modeling are not relevant to a static simulation scenario. The time-invariant CIR may be formulated as

$$h(\tau) = \sum_{i=1}^L \alpha_i \delta(\tau - \tau_i), \quad (14)$$

where each of L taps represents a multipath component of polarity sign-extended (\pm) real amplitude α , multiplied by time delayed Dirac-Delta function. Knowing (14) one can calculate a received signal as a convolution of the CIR with the transmitted signal, i.e., the received signal is a sum of attenuated, delayed and possibly overlapping versions of the transmitted signal, plus white Gaussian noise.

Computing $h(\tau)$ by ray tracing is not as straightforward as one would expect because ray paths, amplitudes and propagation delays are functions of frequency. Ray-tracing evaluates signal at a single frequency, thus computing only narrowband CIR valid for a bandwidth limited transmission. Sub-band divided ray tracing has been proposed in [26], [27] to fully evaluate (14). The method involves multiple simulation runs for a number of center-frequencies of complementary sub-bands. These narrowband CIRs are then combined together in frequency space. Inverse Fourier transform gives the final CIR over wider bandwidth. The amount of computation resources grows linearly with the subdivision granularity, making space and time optimization of simulation tools as proposed here even more important. In the following, only simulations of individual narrowband CIR are presented and analyzed.

Our in-house radio frequency ray-tracing tool has been extended with the concepts outlined in this paper. The tool is configurable with either recursive or non-recursive icosahedral grids for ray launching. It supports all the surveyed double counting avoidance algorithms and new highly configurable Bloom filtering. Distributed wavefronts, as well as plain double and zero counting algorithms have been included as well. The tool has been proven in several projects with telecommunication industry. It is highly optimized GPU-based




Fig. 4. Analytical plane earth loss is compared to simulated values using double counting avoidance by Bloom filtering at 900 MHz, with transmitter and receivers placed 14.6 m and 1.6 m above the ground, respectively.




Fig. 5. Double counting error in the plane earth scenario shows as frequently underestimated path loss, subject to the reception sphere alignment with rays' spatial distribution.

ray tracer using the NVIDIA OptiX ray-tracing engine, which is adapted to the radio frequency simulations. Scene objects are kept in a bounding volume hierarchy entirely on a GPU, with rays generated and traced through the scene in parallel threads. The tool was selected over the commercially available ray tracers because it is fully customizable at source code level. Knowing all the intricate details of the implementation, shortcuts, such as dividing surfaces into tiles and other closely-guarded trade secrets of commercial tools, could not bias the evaluation.

Before focusing on the CIR analysis we demonstrate the effect of double counting errors on the fundamental plane earth propagation scenario by using our ray-tracing tool. The transmission point was placed 14.6 m above the ground whereas the reception points were assumed to be at the height of 1.6 m above the ground, which is consistent with the outdoor scenario presented later. 4 km by 4 km ground plane having electric properties of an asphalt mixture ($\epsilon_r = 5.72, \mu_r = 1, \sigma = 0.0005$ S/m) was the only reflecting obstacle. In Fig. 4 plane earth loss using Bloom filter double counting avoidance with low false positives rate ($p = 10^{-4}$) is plotted along with the analytical plane earth loss reference [28]

$$\frac{P_r}{P_t} = 2 \left(\frac{\lambda}{4\pi r} \right)^2 \left[1 - \cos k \frac{2h_m h_b}{r} \right] \quad (15)$$

The simulated path loss matches the reference for $r > 70$ m, with the expected 30 dB per decade loss above 200 m. The disagreement for shorter distances is normal because (15) assumes the transmission and reception point heights to be small compared to the ground distance r .

Double counting error in simulated plane earth propagation scenario is demonstrated in Fig. 5. Periodic misalignment of reception spheres with the spatial distribution of rays leads to frequent underestimation of the path loss. On the other hand, the ZC algorithm may occasionally completely miss a wavefront, which shows as a sporadic loss overestimate.

More complex indoor and outdoor scenarios were examined in the following evaluation runs. Narrowband channel impulse responses were computed at several observation points with different double counting avoidance and ray launching algorithms in place. The geometry, the transmission sources and the observation points for the two scenarios are presented in Fig. 6 and Fig. 7. The indoor scenario extends over a single floor. The internal walls are made of heavy plasterboard ($\epsilon_r = 6, \mu_r = 1, \sigma = 0.05$ S/m) The outer walls, ground and ceiling material is steel-reinforced concrete ($\epsilon_r = 9, \mu_r = 1, \sigma = 0.09$ S/m). Most of the buildings in the outdoor scenario are made from steel-reinforced concrete, whereas the ground plane is assumed to have asphalt coating. Four sample signal paths are illustrated per scenario with thousands more actually simulated. The observation point in the indoor scenario was placed 0.85 m above the ground in the corner office with the transmission point in the main corridor 15 cm higher than the observation point. In the urban scenario, the transmission point was placed on a building rooftop 14.6 m above the ground and the reception point 1.6 m above the ground in the vicinity of some high-rise buildings, which are the heights introduced earlier in the plane earth scenario. Straight Euclidean paths between the transmission source and the reception points in the presented setups were 9.2 m and 104.3 m in length. Both, the transmission and the reception were simulated assuming ideal vertically oriented dipole antennas. Indoor simulations were run at 4 GHz center frequency of the impulse radio that we use in the research of UWB localization techniques, whereas the urban simulation took place at 900 MHz GSM (Global System for Mobile Communications) frequency. If one employs dipole-like UWB omnidirectional antenna such as Cobham XPO2V-2.0-18.00/1397 in indoor scenario, the Fraunhofer distance valid for free space evaluates at $2D^2/\lambda = 2 \cdot 10.4^2 / 7.49 \approx 29$ cm. The boundary of the radiative near-field region in realistic scenario can be significantly larger. For instance, assuming that the largest dimension of the GSM antenna in outdoor scenario is 2 m, e.g. Kathrein 742 265, the




Fig. 6. Indoor scenario without a ceiling showing the locations of the observation point (blue dot) and the location of the transmitter (red dot) Four sample ray paths involving multiple reflections and transmissions are outlined in red.




Fig. 7. Outdoor urban scenario used in the double counting avoidance evaluation, including four sample ray paths between the transmission point (red dot) and the observation point (blue dot) The simulated paths consisted of reflections and edge diffractions.

far-field region would span several kilometers [29]. However, brute force 3D ray tracing is fully capable of predicting fields within the Fresnel region because the collective interactions of all or most of the scatterers are systematically considered [30].

Multipath signal components were limited by simulation depth and by the individual path signal loss. The allowed ray-tracing depth in indoor scenario was 30 consecutive ray interactions, either reflections or refractions. Note that indoor geometry has low number of potential diffraction edges, thus a diffraction phenomenon was not accounted for in indoor scenario. However, diffraction contribution to the received signal in urban scenario is significant, whereas signal transmission through steel-concrete buildings is usually omitted. Being an open space, 20 consecutive reflections or single-edge

diffractions were sufficient in outdoor scenario to arrive at stable CIR. Up to 150 dB signal loss per multipath component was allowed in both scenarios.

In indoor scenario 167.772.162 rays were initially launched per simulation run, which typically took 9 hours to complete. Fewer interactions in urban scenario allowed 4 times more rays per simulation in shorter 6-hour time frame. Template icosahedral grids with matching number of points are recursive grids at the refinement level of 12 and 13, whereas non-recursive grids should be configured with 4096 and 8192 partitioning sections. The criterion in a selection of the above number of rays was the maximum indoor reception sphere radius of 15 cm for the theoretically longest ray path, given the scene dimensions and the allowed number of interactions. The maximum reception sphere radius in the urban scenario was set to 60 cm. Actual reception sphere radii for most simulated paths are significantly smaller. For example, taking that maximum angular distance of the recursive grid at level 12 is approximately 3.23×10^{-4} radians, the reception sphere radius for rays travelling 100 m is only 1.87 cm.

For each scenario and both ray launching algorithms the double counting avoidance by CS was first simulated. The exact number of different wavefronts n was then used for setting up Bloom filter parameters. As suggested previously, one can get approximate number of different wavefronts by discounting the total number of reception sphere hits if CS is not available. Table I gives the filter parameters for each scenario and both ray launching algorithms. Four false positive rates were selected as a geometric progression with common ratio 10, starting with $p = 0.1$. The optimal number of hash functions is 4 at the highest probability rate up to 14 at the lowest probability rate. Note that the optimal k depends only on p as (7) can be simplified to $-\log_2 p$. The filter sizes shown in the table are rounded up to 32-bit boundaries for better memory alignment. Memory requirements range from a few up to several tens of Kbytes. Outdoor scenario is generally less memory demanding with an order of magnitude less wavefronts hitting the observation point. Larger reception sphere radii due to a non-recursive grid construction in the ray launching step led to a larger number of detected wavefronts and demanded approximately 1.7% and 12.9% increase in the filter sizes for indoor and outdoor scenarios, respectively.

Considering channel impulse response quality, a double counting error shows itself as an extra tap, overlapping or in close temporal proximity to the legitimate tap. On the other hand, false positives, either in the context of Bloom filters or as invalid decisions by the AI or AD, exhibit as missing taps with respect to the exact solution. Table II shows tap counts for algorithms under consideration for the two scenarios. Recursive and non-recursive ray-launching grids are presented separately. The output of CS is regarded to be the reference solution against which the quality of other approaches is measured.

The fact that DC and DW allow all rays to affect the CIR shows as the highest error in the number of taps L . Well-known theoretical 20.9% double counting rate effectively doubles to a value above 41% in the indoor scenario when maximum angle of icosahedral grids is used and the refraction-

TABLE I
BLOOM FILTER PARAMETERS

p	$[k_{\text{opt}}]$	Indoor, $n = 12, 105$		Outdoor, $n = 1, 809$	
		m^1	Bytes	m^1	Bytes
Recursive icosahedral grids					
0.0001	14	232,064	29,008	34,688	4,336
0.001	10	174,048	21,756	26,016	3,252
0.01	7	116,032	14,504	17,344	2,168
0.1	4	58,016	7,252	8,672	1,084
p	$[k_{\text{opt}}]$	Indoor, $n = 12, 309$		Outdoor, $n = 2, 037$	
		m^1	Bytes	m^1	Bytes
Non-recursive icosahedral grids					
0.0001	14	235,968	29,496	39,072	4,884
0.001	10	176,992	22,124	29,312	3,664
0.01	7	117,984	14,748	19,552	2,444
0.1	4	59,008	7,376	9,792	1,224

¹ rounded up to 32-bit boundary

related ray-spacing artefacts occur. Significantly larger error of at least 300% is reported in the outdoor scenario due to the highly irregular spacing of rays on the diffraction-modelling Keller cones. Note that we do not attempt to correct the reception sphere radii for the closer origin of diffraction rays as explained in Section II. On the other hand, ZC allows some misses while introducing new auxiliary rays to minimize the uncovered area. Because rays cannot be generated uniformly and maximum icosahedral grid angle is used in determining reception sphere radii, double count events still occur. Further, the decreased and irregular ray spacing also contribute to double counts. Thus, in our experiments even ZC produced a surplus of CIR taps above 10% without and 180% with the diffraction phenomenon included. Here one must be aware that the number of erroneous double taps is even larger as some taps are actually missing.

The AI and the AD heuristics both base set membership decisions on the angle similarity. In order to avoid the false negatives, i.e., double counts, the angle between newly arrived ray and rays in a set should differ at least $2/\sqrt{3}\alpha$. However, due to the refraction-decreased ray spacing and irregular spacing of diffraction rays, false negatives are still possible, especially for the diffraction rays. The angle criterion was doubled in simulations to avoid most false negatives in the indoor scenario, whereas in the outdoor scenario false negatives still prevailed. Additional source of decision errors in the AD heuristic are different path lengths of adjacent rays. Paths within 0.1% length difference were assumed equal. Nevertheless, even when angle and distance criteria are used in conjunction, many rays are wrongly associated with wavefronts. AI is equally prone to wrong decisions with a few percentage points smaller error indoors compared to AD, which is still significant, with 21% less taps than expected being the best result in Table II.

Bloom filters configured with low false positive rates expectedly recognize all or most of the taps. In the indoor scenario there was a single missed tap at $p = 10^{-4}$. The configured

rate would allow 1 or 2 errors in average at fully populated set. In the urban scenario the first missing taps occurred at $p = 10^{-2}$ as less wavefronts hit the reception sphere, which allows optimal performance using shorter filters. Increasing p gradually corrupts channel impulse responses. However, even at marginal 10^{-1} Bloom filtering performs significantly better than any sub-optimal double counting avoidance. Further, the percentage of missed taps grows slower than the configured false positive rate. This is expected behavior. In order to maintain a fixed average rate the filter size should be growing linearly with the number of inserted elements, whereas the proposed double counting avoidance uses full size filters from the start.

The non-recursive icosahedral grids generally perform worse than the recursive grids as a ray-launching template in terms of double and missed taps. This can be explained by less uniform distribution of non-recursively generated rays with larger maximum and smaller minimum angular distances at the comparable number of grid points. Optimal wavefront differentiation produces more valid taps because larger reception spheres are used. Angle based AD and AI are more error prone, whereas Bloom filtering—being independent of angular distribution—performs equally well on both grid templates.

Additional columns in Table II give a continuous signal loss and its relative error to the reference CS loss, calculated from the complex tap amplitudes. Further, having multiple observation points, we calculated the path loss exponent for the two scenarios as 2.4 indoor and 2.8 outdoor. The outdoor exponent only seems to be in disagreement with the more common value around 4 that can be found in many empirical propagation models within macrocells. Our loss exponent is primarily affected by the scenario size, which is 400 m by 400 m, whereas the empirical models provide a good approximation at distances larger than the break point, typically several hundred meters from the transmission source. Due to short distances the direct and the reflected path lengths are not similar and highly fluctuates. Additionally, simulations assumed vertically polarized transmission, which is prone to significant ground-reflected magnitude variations around the Brewster angle. In the presented case the angle with zero magnitude reflection is at 67 degree incidence (ground index of refraction is 2.39). The simulated scenario is closer to microcell environment. The base station antenna is relatively low with respect to some larger buildings in the area. In addition to free-space propagation, multiple reflections and diffractions around the vertical edges have a significant impact on the signal loss. A better empirical model to compare the presented scenario with would be a dual-slope model in which path loss exponent for short distances is known to be around 2.

Field CIR measurements do not produce individual taps with arbitrary time precision. The actual resolution is a reciprocal of the probing pulse bandwidth [31]. Instead of the individual taps one gets bin amplitudes, which are functionally dependent on those taps. The quality of CIR can then be given as a mean and variance of amplitude errors with respect to the reference CS. In order to analyze the double counting error in that context, power delay profiles with 5 ns resolution (200

TABLE II
CHANNEL IMPULSE RESPONSE QUALITY

	Indoor scenario					Outdoor scenario				
	Taps	Error	%	Loss [dB]	%	Taps	Error	%	Loss [dB]	%
Recursive icosahedral grids										
Characteristic Sequences (CS)	12,105	0	0.00	77.37	0.00	1,809	0	0.00	65.88	0.00
Double Counting (DC)	17,078	4,973	41.08	68.74	-11.15	7,348	5,539	306.19	55.58	-15.63
Zero Counting (ZC)	13,483	1,378	11.38	79.73	3.05	5,070	3,261	180.27	59.49	-9.70
Distributed Wavefronts (DW)	17,078	4,973	41.08	74.76	-3.37	7,348	5,539	306.19	55.54	-15.70
Angle and Interactions (AI)	9,550	-2,555	-21.11	77.05	-0.41	5,036	3,227	178.39	58.82	-10.72
Angle and Distances (AD)	8,728	-3,377	-27.90	76.43	-1.21	4,861	3,052	168.71	58.65	-10.97
Bloom Filter (BF 0.0001)	12,104	-1	-0.01	77.37	0.00	1,809	0	0.00	65.89	0.02
Bloom Filter (BF 0.001)	12,103	-2	-0.02	77.38	0.01	1,809	0	0.00	65.86	-0.03
Bloom Filter (BF 0.01)	12,089	-16	-0.13	77.32	-0.06	1,808	-1	-0.06	66.04	0.24
Bloom Filter (BF 0.1)	11,731	-374	-3.09	76.01	-1.76	1,753	-56	-3.10	66.11	0.35
Non-recursive icosahedral grids										
Characteristic Sequences (CS)	12,309	0	0.00	76.16	0.00	2,037	0	0.00	66.81	0.00
Double Counting (DC)	18,476	6,167	50.10	62.53	-17.90	8,254	6,217	305.20	62.60	-6.30
Zero Counting (ZC)	14,424	2,115	17.18	64.64	-15.13	5,710	3,673	180.31	61.16	-8.46
Distributed Wavefronts (DW)	18,476	6,167	50.10	68.16	-10.50	8,254	6,217	305.20	60.90	-8.85
Angle and Interactions (AI)	9,657	-2,652	-21.55	75.86	-0.39	5,606	3,569	175.21	64.53	-3.41
Angle and Distances (AD)	8,845	-3,464	-28.14	75.87	-0.38	5,366	3,329	163.43	63.95	-4.28
Bloom Filter (BF 0.0001)	12,309	0	0.00	76.12	-0.05	2,037	0	0.00	66.74	-0.10
Bloom Filter (BF 0.001)	12,308	-1	-0.01	76.13	-0.04	2,037	0	0.00	66.81	0.00
Bloom Filter (BF 0.01)	12,286	-23	-0.19	76.20	0.05	2,035	-2	-0.10	66.79	-0.03
Bloom Filter (BF 0.1)	11,968	-341	-2.77	76.16	0.00	1,980	-57	-2.80	67.58	1.15

MHz bandwidth-limited pulse) were calculated and compared to one another. In Fig. 8 DC, ZC, DW and AI are compared to the optimal CS for the indoor observation point and recursive ray launching. The mean error and its standard deviation are high for both DC and ZC with no double counting avoidance in place. The numerous extra taps show as positive error bias. DW manages to reduce the mean error, but still produces unacceptably high variations for good CIR estimate. AI and AD are only marginally better than CW with slightly less amplitude variations. On the other hand, Bloom power delay profiles in Fig. 9 show far superior matches, with the more stringent variants generating practically optimal results. In the urban scenario higher deviations are observed for non-Bloom heuristics but with similar conclusions. Full numerical results are given in Table III.

The loss metric is far less informative about the algorithmic errors due to its cumulative nature. Larger tap amplitudes overshadow smaller, while the time-related CIR information content is completely lost. Nevertheless, we provide comparison of measured path loss against simulations considering different approaches to double counting elimination. A measurement campaign has been launched in the city of Ljubljana in close vicinity of the outdoor scenario presented earlier. In cooperation with Telecom Slovenia we measured path loss for a number of points at 2m above the ground within UMTS (Universal Mobile Telecommunications System)

cell. The locations of measurements are shown in Fig. 10. The measurement equipment consisted of the Ascom TEMS Investigation (Infovista as of 2016) air interface test tool connected to the Sony Ericsson K800i mobile terminal. The external antenna was mounted on a vehicle rooftop. Common pilot channel (CPICH) received signal code powers (RSCP) of the serving cell expressed in dBm were captured together with the GPS location data. The measured loss in Fig. 10 is a difference between the 27.7 dBm transmission power and the reported measurements. Kathrein 742 265 antenna radiation pattern with 4 degrees electrical tilt and 16.15 dBd gain was simulated to match on-field conditions as closely as possible. In order to account for the unknown equipment loss a calibration factor with respect to the CS algorithm was computed and applied to all simulated values. Measured and simulated losses are presented in Fig. 11. In contrast to the plane earth loss, values presented here are total losses, including directional antenna gains and equipment losses. The quality of ray-tracing estimates is then evaluated in terms of standard deviation of the difference between measured and simulated path loss. Table IV lists the estimated values. The BF with the lowest false positives rate practically halves the error of sub-optimal double counting approaches introduced by the other algorithms. However, the standard deviation of 5 dB is still high, which can be attributed to the modelling errors. The area is populated with trees that are not accounted

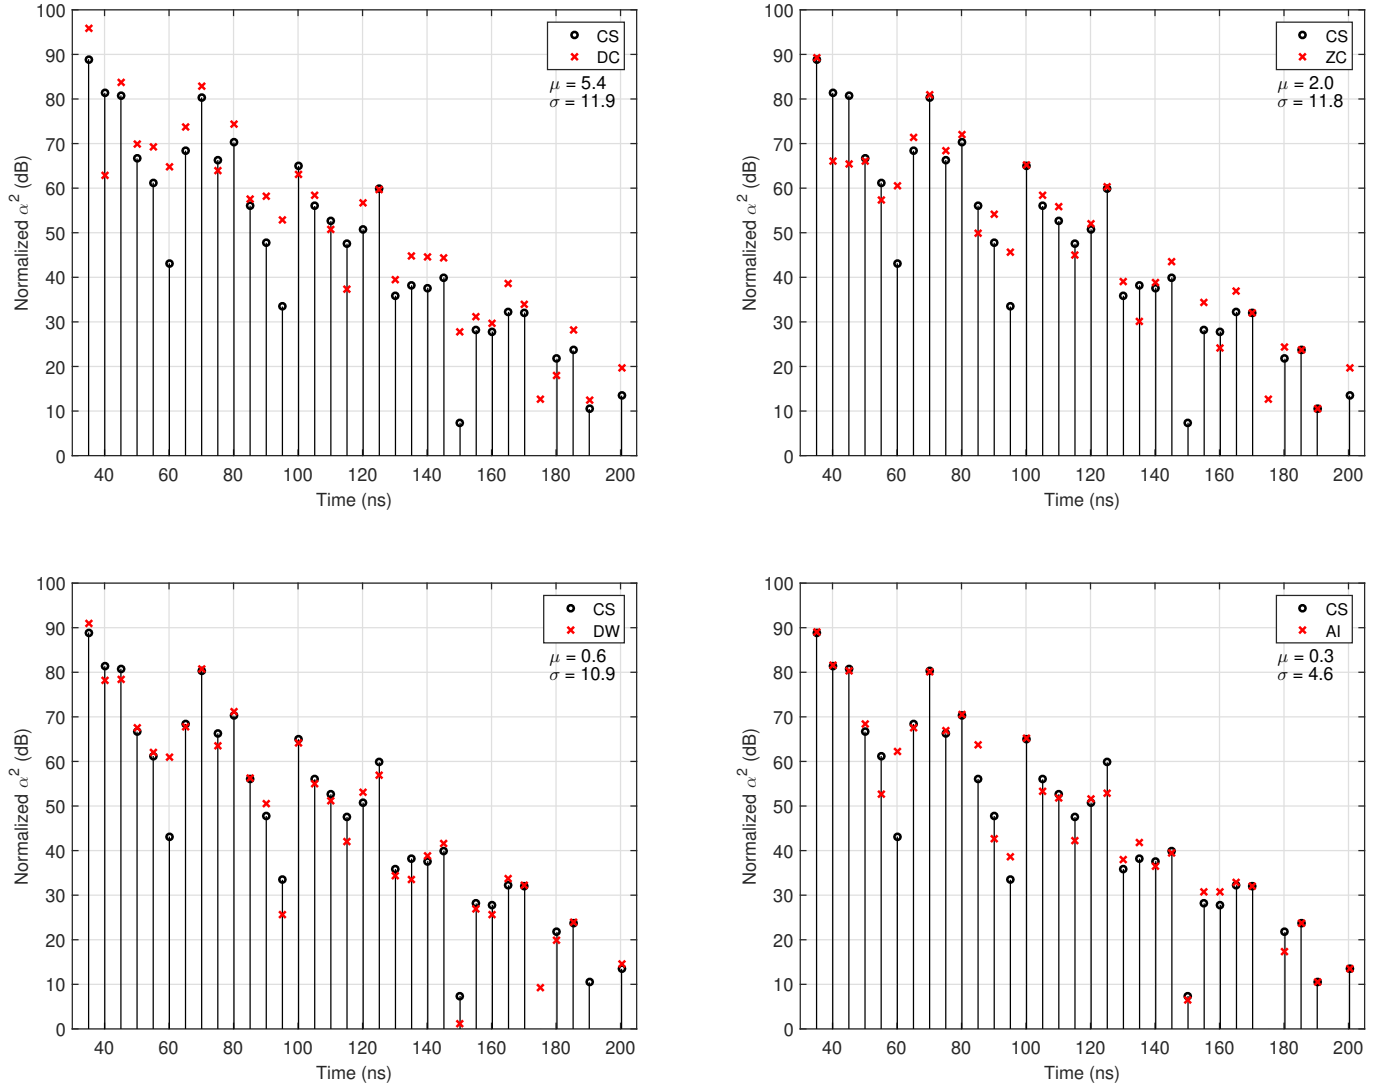


Fig. 8. Accuracy of simulated power delay profiles with 5 ns resolution if double counting is allowed or less efficient avoidance heuristics are used (indoor scenario and recursive grids). Standard deviation of the amplitude errors is between 4 and 12 dB. DC and ZC are additionally penalized by positive mean offset of several dB.

TABLE III
BINNED PROFILE VARIATIONS

	Recursive icosahedral grids				Non-recursive icosahedral grids			
	Indoor		Outdoor		Indoor		Outdoor	
	μ	σ	μ	σ	μ	σ	μ	σ
Double Counting (DC)	5.45	11.86	10.95	5.73	4.62	7.27	9.60	11.77
Zero Counting (ZC)	2.04	11.80	7.47	5.57	1.42	8.64	7.41	6.32
Distributed Wavefronts (DW)	0.58	10.86	10.87	5.90	0.10	6.75	11.05	5.65
Angle and Interactions (AI)	0.31	4.64	7.27	5.24	-0.33	4.24	7.88	5.17
Angle and Distances (AD)	-0.22	7.05	7.49	5.14	-0.92	6.88	7.51	5.03
Bloom Filter (BF 0.0001)	0.01	0.06	0.00	0.00	0.00	0.01	0.00	0.00
Bloom Filter (BF 0.001)	-0.03	0.18	-0.00	0.00	0.00	0.02	-0.00	0.00
Bloom Filter (BF 0.01)	-0.06	0.51	-0.00	0.00	-0.20	0.87	-0.00	0.01
Bloom Filter (BF 0.1)	2.69	11.59	-0.14	0.23	-1.23	4.75	-0.36	0.75

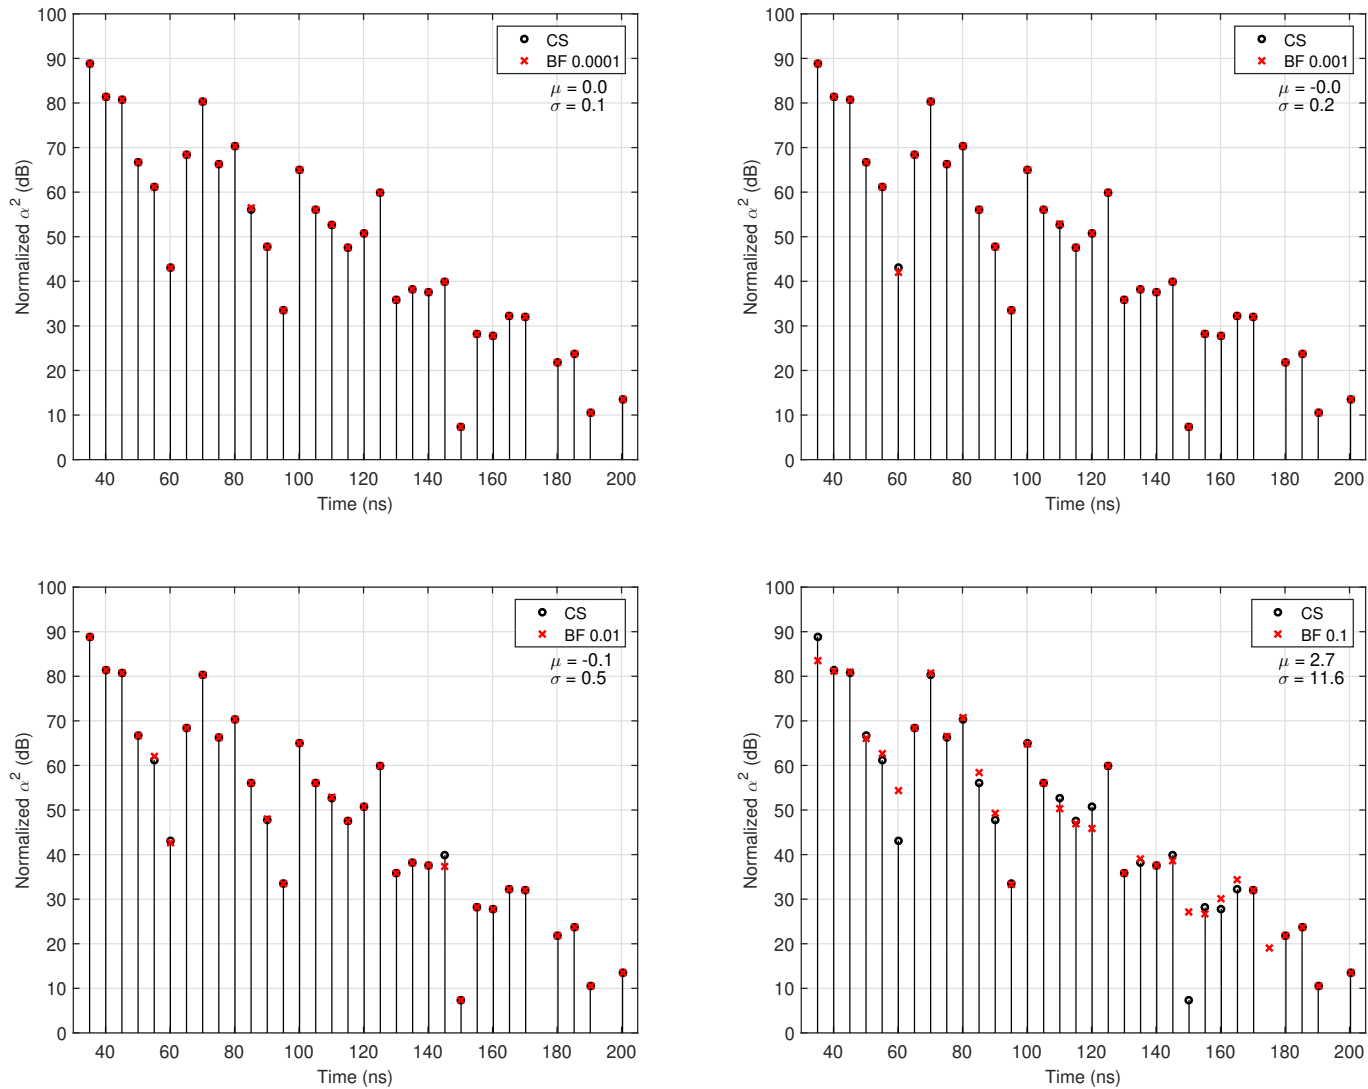


Fig. 9. Accuracy of simulated power delay profiles with 5 ns resolution when Bloom filters are used (indoor scenario and recursive grids). Nearly optimal matches are produced for the lowest false positive rates. Moreover, even the highest rate of 0.1 is better than any non-Bloom suboptimal approach.

TABLE IV
SIMULATED VS. MEASURED PATH LOSS [dB]

	μ	σ
Bloom Filter (BF 0.0001)	0.0	5.2
Double Counting (DC)	10.7	10.4
Zero Counting (ZC)	7.6	10.0
Distributed Wavefronts (DW)	10.7	10.4
Angle and Interactions (AI)	10.7	10.4
Angle and Distances (AD)	10.7	10.4

for in our simulations. The attenuation of rays propagating through the foliage contributes to the observed deviations due to multiple scattering phenomena.

Excessive memory storage requirements are the main reason why optimal CS cannot be used in practice when impulse responses at several observation points need to be estimated.

Running a simulation per single observation point is usually not an option as each run can easily take several hours to complete. Bloom filters are of the same $\mathcal{O}(n)$ space bound as AI and AD, but with significantly smaller constant factor while providing practically equivalent solutions to $\mathcal{O}(n \times depth)$ space bounded CS. Actual memory requirements for the two scenarios while using recursive icosahedral grids are presented in Table V. The incident angles memorized by AD and AI are represented as direction vectors consisting of 3 floats, distances as single floats and the number of interactions as single byte values. For a given value of *depth*, the ratio to the CS memory size remains constant across experiments as shown in Fig. 12. The suggested indoor Bloom filters at $p = 10^{-4}$ take 4% of the CS space, whereas the suggested outdoor filtering at $p = 10^{-2}$ consumes only 3% of the CS space. Both configurations produce CIRs that marginally differ from the ones obtained by the optimal CS.

Further advantage of Bloom filtering for a double counting




Fig. 10. Measurement points in the outdoor path loss evaluation experiment (red dots). The location of transmitter (blue point) is N46.037832°, E14.497457° in the city of Ljubljana.




Fig. 11. Comparison of path loss measurements and predictions along the measurement route within 2.1 GHz UMTS cell

TABLE V
DOUBLE COUNTING AVOIDANCE MEMORY REQUIREMENTS

	Recursive icosahedral grids			
	Indoor		Outdoor	
	Bytes	%	Bytes	%
Characteristic Sequences (CS)	726,300	100.0	72,360	100.0
Angle and Interactions (AI)	157,365	21.7	23,517	32.5
Angle and Distances (AD)	193,680	26.7	28,944	40.0
Bloom Filter (BF 0.0001)	29,008	4.0	4,336	6.0
Bloom Filter (BF 0.001)	21,756	3.0	3,252	4.5
Bloom Filter (BF 0.01)	14,504	2.0	2,168	3.0
Bloom Filter (BF 0.1)	7,252	1.0	1,084	1.5




Fig. 12. Relative amount of memory required by the double hit avoidance algorithms with respect to the optimal CS while simulating 30 indoor interactions per ray, given in descending order

avoidance over the CS, AD and AI algorithms are constant membership query times. However, the ray-tracing time far surpasses the total membership query time for a single observation point. Hence no significant difference has been observed in the running times of DC, DW, AI, AD and BFs. This was not the case for the CS membership queries, which involve longer characteristic sequences. The total running times of the CS runs increased by factor 1.2 in average. Finally, smaller reception sphere radius and auxiliary rays had significant impact on the ZC simulations, with on average of 2.9-times longer execution than the best performing algorithms.

VI. DISCUSSION

We use least complex variant of Bloom filter for the double counting avoidance problem. In Section III we briefly mention the possibility of using scalable, dynamic and split filters to address the unknown set cardinality. Other types of filters may also prove beneficial.

False positive rate can be reduced by allowing some false negatives, i.e., double counts. There is less or no chance of a double hit if ray passes the reception sphere sufficiently close to the sphere central point due to the use of regular grids in ray launching process. Such rays may be simply omitted from the set. The criteria are not straightforward, mainly because of the reduced ray spacing. The other option is the use of filters that support probabilistic deletions. Retouched Bloom filter [32] randomly resets some table bits to trade off false positives against false negatives. More complex selective clearing process based on training sequences can target only troublesome false positives, such as those affecting dominant propagation paths.

Counting Bloom filter keeps a small counter for each table bit [33], [34]. Instead of setting a bit, corresponding counter is incremented. Set element removal is accomplished by decrementing the matching counters. One possible use of counting filter would be to delete recognized double hits from

the table and hopefully reduce the overall false positive rate; or to maintain equivalent rate with shorter tables. The rationale is that more than two hits per wavefront are rare, thus the deletion is more likely to reduce the positives rate than to introduce false negatives.

Probabilistic deletions that are not restricted to double hits without introducing any false negatives are also achievable as proposed in DIBF [35]. The method involves additional record keeping.

Stable Bloom filter [36] has been designed specifically to address duplicate detection. The filter introduces both types of errors, but with improved rates in the context of data stream processing. The problem shows some similarities with the double counting avoidance. We leave further improvements, including exploration of the above concepts, for future work.

VII. CONCLUSION

This paper proposed near-optimal double counting avoidance of wavefronts by Bloom filtering that can be applied to indoor and outdoor scenarios. Even when configured with the false positive rate as high as 10^{-4} , simulation space reduction proves to be considerable. Further, the choice of a ray launching template should not be underestimated. Proper analytical bounds of sphere radii are provided to capture at least one ray per passing wavefront. The paper showed that channel impulse response simulations free of most algorithmic errors are possible in a single simulation run for the number of observation points on the scale of more conventional coverage predictions. The presented ray-tracing optimization is particularly important for computing architectures where memory space is at premium.

ACKNOWLEDGMENT

The author would like to thank the colleagues at Telekom Slovenije, d.d. for carrying out UMTS signal measurement campaign.

REFERENCES

- [1] F. Ikegami, T. Takeuchi, and S. Yoshida, "Theoretical prediction of mean field strength for urban mobile radio," *IEEE Trans. Antennas Propag.*, vol. 39, no. 3, pp. 299–302, Mar. 1991.
- [2] G. Durgin, N. Patwari, and T. Rappaport, "An advanced 3D ray launching method for wireless propagation prediction," in *Proc. of the 47th IEEE Vehicular Technology Conference*, vol. 2, Phoenix, AZ, 1997, pp. 785–789.
- [3] N. Noori, A. Shishegar, and E. Jedari, "A new double counting cancellation technique for three-dimensional ray launching method," in *IEEE Antennas and Propagation Society International Symposium*, Albuquerque, NM, 2006, pp. 2185–2188.
- [4] Z. Chen, H. L. Bertoni, and A. Delis, "Progressive and approximate techniques in ray-tracing-based radio wave propagation prediction models," *IEEE Trans. Antennas Propag.*, vol. 52, no. 1, pp. 240–251, 2004.
- [5] Y. Tao, H. Lin, and H. Bao, "GPU-based shooting and bouncing ray method for fast RCS prediction," *IEEE Trans. Antennas Propag.*, vol. 58, no. 2, pp. 494–502, 2010.
- [6] H. Suzuki and A. Mohan, "Ray tube tracing method for predicting indoor channel characteristics map," *Electron. Lett.*, vol. 33, no. 17, pp. 1495–1496, 1997.
- [7] A. Schmitz, T. Rick, T. Karolski, T. Kuhlen, and L. Kobbelt, "Efficient rasterization for outdoor radio wave propagation," *IEEE Trans. Vis. Comput. Graphics*, vol. 17, no. 2, pp. 159–170, 2011.
- [8] J. W. McKown and R. L. Hamilton, "Ray tracing as a design tool for radio networks," *Netwrk. Mag. of Global Internetwork.*, vol. 5, no. 6, pp. 27–30, 1991.
- [9] W. Lu and K. Chan, "Advanced 3D ray tracing method for indoor propagation prediction," *Electron. Lett.*, vol. 34, no. 12, pp. 1259–1260, 1998.
- [10] S. Flores, L. Mayorgas, and F. Jimenez, "Reception algorithms for ray launching modeling of indoor propagation," in *Proc. of IEEE Radio and Wireless Conference*, Colorado Springs, CO, 1998, pp. 261–264.
- [11] Z. Yun, M. Iskander, and Z. Zhang, "Development of a new shooting-and-bouncing ray (SBR) tracing method that avoids ray double counting," in *IEEE Antennas and Propagation Society International Symposium*, vol. 1, 2001, pp. 464–467.
- [12] N. Wang and J.-L. Lee, "Geometric properties of the icosahedral-hexagonal grid on the two-sphere," *SIAM Journal on Scientific Computing*, vol. 33, no. 5, pp. 2536–2559, 2011.
- [13] B. H. Bloom, "Space/time trade-offs in hash coding with allowable errors," *Commun. ACM*, vol. 13, no. 7, pp. 422–426, Jul. 1970.
- [14] A. Broder and M. Mitzenmacher, "Network applications of Bloom filters: A survey," *Internet Math.*, vol. 1, no. 4, pp. 636–646, 2002.
- [15] F. Chang, J. Dean, S. Ghemawat, W. C. Hsieh, D. A. Wallach, M. Burrows, T. Chandra, A. Fikes, and R. E. Gruber, "Bigtable: A distributed storage system for structured data," *ACM Trans. Comput. Syst.*, vol. 26, no. 2, pp. 4:1–4:26, Jun. 2008.
- [16] S. Tarkoma, C. E. Rothenberg, and E. Lagerstedt, "Theory and practice of Bloom filters for distributed systems," *IEEE Commun. Surveys Tuts.*, vol. 14, no. 1, pp. 131–155, 2012.
- [17] J. B. Keller, "Geometrical theory of diffraction," *J. Opt. Soc. Am.*, vol. 52, no. 2, pp. 116–130, 1962.
- [18] P. Bose, H. Guo, E. Kranakis, A. Maheshwari, P. Morin, J. Morrison, M. Smid, and Y. Tang, "On the false-positive rate of Bloom filters," *Information Processing Letters*, vol. 108, no. 4, pp. 210–213, 2008.
- [19] P. S. Almeida, C. Baquero, N. Preguica, and D. Hutchison, "Scalable Bloom filters," *Information Processing Letters*, vol. 101, no. 6, pp. 255–261, 2007.
- [20] D. Guo, J. Wu, H. Chen, Y. Yuan, and X. Luo, "The dynamic Bloom filters," *IEEE Trans. Knowl. Data Eng.*, vol. 22, no. 1, pp. 120–133, Jan. 2010.
- [21] M.-Z. Xiao, Y.-F. Dai, and X.-M. Li, "Split Bloom filter," *Acta Electronica Sinica*, vol. 32, no. 2, pp. 241–245, Feb. 2004.
- [22] M. V. Ramakrishna, "Practical performance of Bloom filters and parallel free-text searching," *Commun. ACM*, vol. 32, no. 10, pp. 1237–1239, Oct. 1989.
- [23] P. L'Ecuyer and P. Hellekalek, "Random and quasi-random point sets," in *Random number generators: selection criteria and testing, Lecture Notes in Statistics*, P. Hellekalek and G. Larcher, Eds., vol. 138. New York, NY: Springer, 1998, pp. 223–265.
- [24] T. M. Wang, P. H. Tseng, Y. C. Chan, and D. B. Lin, "A ray-tracing based fingerprinting for indoor positioning," in *IEEE Com. Soc. Proc.*, Aug. 2015, pp. 1859–1863.
- [25] W. J. Chang and J. H. Tarng, "Effects of bandwidth on observable multipath clustering in outdoor/indoor environments for broadband and ultrawideband wireless systems," *IEEE Trans. Veh. Technol.*, vol. 56, no. 4, pp. 1913–1923, 2007.
- [26] H. Sugahara, Y. Watanabe, T. Ono, K. Okanou, and S. Yamazaki, "Development and experimental evaluations of "RS-2000" - a propagation simulator for UWB systems," in *International Workshop on Ultra Wideband Systems (IEEE Cat. No.04EX812)*, May 2004, pp. 76–80.
- [27] P. Meissner, M. Gan, F. Mani, E. Leitinger, M. Fröhle, C. Oestges, T. Zemen, and K. Witrisal, "On the use of ray tracing for performance prediction of UWB indoor localization systems," in *Proc. of the IEEE International Conference on Communications (ICC'13)*, Jun. 2013, pp. 68–73.
- [28] S. Saunders and A. Aragón-Zavala, *Antennas and Propagation for Wireless Communication Systems*, 2nd ed. Wiley, May 2007.
- [29] A. De, T. K. Sarkar, and M. Salazar-Palma, "Characterization of the far-field environment of antennas located over a ground plane and implications for cellular communication systems," *IEEE Antennas Propag. Mag.*, vol. 52, no. 6, pp. 19–40, 2010.
- [30] Z. Yun and M. F. Iskander, "Ray tracing for radio propagation modeling: Principles and applications," *IEEE Access*, vol. 3, pp. 1089–1100, 2015.
- [31] H. Hashemi, "Impulse response modeling of indoor radio propagation channels," *IEEE J. Sel. Areas Commun.*, vol. 11, no. 7, pp. 967–978, Sep. 1993.
- [32] B. Donnet, B. Baynat, and T. Friedman, "Retouched Bloom filters: Allowing networked applications to trade off selected false positives

against false negatives,” in *Proc. of the 2006 ACM CoNEXT Conference*, New York, NY, 2006, pp. 13:1–13:12.

- [33] K.-Y. Whang, B. T. Vander-Zanden, and H. M. Taylor, “A linear-time probabilistic counting algorithm for database applications,” *ACM Trans. Database Syst.*, vol. 15, no. 2, pp. 208–229, Jun. 1990.
- [34] D. Ficara, S. Giordano, G. Procissi, and F. Vitucci, “Multilayer compressed counting Bloom filters,” in *IEEE INFOCOM 2008 - The 27th Conference on Computer Communications*, Apr. 2008, pp. 311–315.
- [35] C. E. Rothenberg, C. A. B. Macapuna, F. L. Verdi, and M. F. Magalhaes, “The deletable Bloom filter: a new member of the Bloom family,” *IEEE Commun. Lett.*, vol. 14, no. 6, pp. 557–559, Jun. 2010.
- [36] F. Deng and D. Rafiei, “Approximately detecting duplicates for streaming data using stable Bloom filters,” in *Proc. of the 2006 ACM SIGMOD International Conference on Management of Data*, New York, NY, 2006, pp. 25–36.



Roman Novak received his MSc and PhD degrees in computer science from the University of Ljubljana in 1995 and 1998, respectively. He is a researcher at the Department of Communication Systems of the Jožef Stefan Institute in Ljubljana, Slovenia since 1992 and an assistant professor at the Jožef Stefan International Postgraduate School since 2007. He was a visiting scientist at the Joanneum Research Forschungsgesellschaft, Graz, Austria in 2008. He participated in several major projects in the field of design and development of telecommunication systems. His latest research interests include radiofrequency ray tracing and parallel algorithms in general. His previous research interests have been focused on routing in communication networks, information security topics, as well as on signal processing for nuclear spectroscopy and others. He is a member of the IEEE and the ACM.