Quantifying Variability in Silicon Photonics: A Stochastic Study of Sidewall Roughness and Its Impact on Waveguide Performance

Naveen Kumar[a], Natale G. Pruiti[b], Jeremi Januszewicz[c], Eugenio Di Gaetano[d], Luiz Felipe Aguinsky[e], Marc Sorel[f], Douglas J Paul[g], Kevin Gallacher[h], Vihar Georgiev[i]
James Watt School of Engineering, *University of Glasgow,* Glasgow, UK
Email: [a]naveen.kumar@glasgow.ac.uk, [b]natale.pruiti@glasgow.ac.uk, [c]jeremi.januszewicz@glasgow.ac.uk, [d]eugenio.digaetano@glasgow.ac.uk, [e]luiz.aguinsky@glasgow.ac.uk, [f]marc.sorel@glasgow.ac.uk, [g]douglas.paul@glasgow.ac.uk, [h]kevin.gallacher@glasgow.ac.uk, [i]vihar.georgiev@glasgow.ac.uk

***Abstract*— This work presents a comprehensive simulation-based study of surface roughness-induced loss and phase variation in silicon photonic waveguides extended to centimeter-scale lengths. Our results indicate that subtle changes in the surface roughness profile can yield measurable variations in overall device performance, despite maintaining the same root-mean-square (RMS) roughness amplitude and correlation length. We compare waveguides' performance as a function of different roughness parameters and use a variational finite-difference time-domain framework to handle large-scale computations. The data reveal that even modest increases in roughness (σ_{RMS}) or correlation length can lead to significant scattering loss and phase uncertainties when segments are stitched together randomly. These findings underscore the need for precise control over the sidewall fabrication process and highlight the value of statistical modeling techniques in designing robust photonic devices.**

Keywords*— *Silicon Waveguide, Roughness, Correlation length, FDTD, VarFDTD, Loss variation, Phase Variation

I. Introduction

Silicon photonics has emerged as a pivotal technology for integrated optical systems, providing high-density, low-cost solutions for data center interconnects, on-chip communication, and a broad array of sensing applications [1][2]. Waveguides fabricated on silicon-on-insulator substrates form the backbone of these systems [3][4], exploiting the high refractive index contrast between silicon and silicon dioxide to tightly confine light within submicrometer geometries. However, as waveguide dimensions shrink, scattering losses caused by nanoscale sidewall roughness become increasingly significant [5]. Even small deviations, on the order of a few nanometers can translate to measurable increases in attenuation and phase mismatch, which can severely impact performance in longer, more complex photonic circuits [6].

Fabrication imperfections typically manifest in the form of random peaks and valleys along the waveguide perimeter [7] [8]. Their characteristics are often described by a root-mean-square (RMS) amplitude and a correlation length that establishes how frequently the boundary profile fluctuates [9]. These roughness parameters are crucial, yet they do not alone capture the complete picture, because the exact spatial placement of roughness features also matters. While a single waveguide segment of a few micrometers may show manageable losses, concatenating many such segments to reach centimeter-scale lengths amplifies the cumulative effects of scattering. Furthermore, practical photonic systems operate under conditions where thousands of individual waveguide segments may need to be assembled or interconnected [10]. Each segment, although nominally similar, has its own unique random roughness "fingerprint," which can produce substantial variations in insertion loss and phase shifts across a large device population [11] [12]. Such variability poses challenges for high-volume manufacturing, where consistent device behavior is essential for widespread adoption.

The work reported here addresses the pressing need to understand and statistically quantify these random variations. By generating an extensive set of waveguide samples with controlled yet stochastically distributed sidewall roughness, we employ a full 3D finite-difference time-domain (FDTD) and variational finite-difference time-domain (VarFDTD) approach to simulate electromagnetic propagation with reduced computational overhead compared to full 3D models. We then create one-centimeter-long waveguides by stitching together random sequences of 10 µm segments drawn from a library of samples, thereby revealing how global device properties emerge from local roughness fluctuations. Our findings underscore the importance of advanced fabrication techniques, such as refined lithography and etch processes, aimed at minimizing roughness in critical photonic infrastructure. This study also highlights the necessity for statistical design margins and reliability analysis in commercial photonic integrated circuit applications.

II. Results and Discussion

In Fig. 1(a), a 3D perspective of a silicon waveguide wrapped in SiO_2 (hidden for clarity) with roughness (σ_{RMS}) of 5 nm and a correlation length (L_C) of 100 nm in the x, y, and z directions. This depiction emphasizes the presence of roughness along the waveguide's top surface and sidewalls, which differs from the ideal, perfectly rectangular geometry that designers typically aim to achieve. Even though the absolute scale of these irregularities is small relative to the wavelength of light used in telecommunications (usually around 1.55 µm), they can still produce a significant scattering effect over longer propagation distances [13]. Fig. 1(b), presented from the top (XY) view, provides a clearer visualization of how lateral roughness modifies the waveguide width from its nominal size [14]. The correlation length of 100 nm means that localized surface variations are not merely random bumps but rather have a spatial coherence over that scale. Fig. 1(c) shows the optical mode profile in the plane of the waveguide cross-section, viewed from the front (across the waveguide cross-section), depicting that the considered simulation region is quite large enough, leading to higher accuracy in the obtained results. To illustrate the baseline scenario, Fig. 1(d) captures the electric field for a perfectly smooth waveguide without roughness. The field contours here

are symmetrical, with minimal radiation extending beyond the core boundaries, except for the expected evanescent decay. By contrast, Fig. 1(e) shows how the electric field evolves in the presence of roughness with σ_{RMS} = 1 nm at L_C = 100 nm. Although 1 nm is often considered a small deviation relative to the overall waveguide dimensions, the figure demonstrates that even these slight imperfections can cause noticeable distortions and scattering in the local field [15]. Fig. 1(f) provides loss variation across the waveguide length for different σ_{RMS}. A higher σ_{RMS} leads to more intense scattering of the guided mode, leading to larger propagation losses over distance [16].

Fig. 2 illustrates the impact of σ_{RMS} and L_C on both loss and phase for a 10 µm long silicon waveguide, comparing cases of roughness on all surfaces versus roughness confined only to the sidewalls (YZ plane) under varying σ_{RMS}. As the σ_{RMS} increases, there is a rise in the propagation loss (Fig. 2(a)), due to a greater fraction of light scattered out of the core [2]. Phase changes also increase (Fig. 2(b)) with greater σ_{RMS} because the effective index perturbation accumulates along the 10 µm waveguide. All-around roughness, given the same σ_{RMS}, leads to higher loss and more phase deviation because the top surface, bottom interface, and sidewalls all introduce scattered fields into the simulation space. Sidewall-only roughness tends to isolate the scattering to the perimeter regions where the field intensity may be largest for a horizontally polarised mode, though the lack of top and bottom roughness can reduce the overall scattering volume. When the L_C is reduced to 10 µm for X & Y axes, for higher values of σ_{RMS} (>6nm), loss is a bit less (Fig. 2(c)) for all-around roughness due to effective volume increases as compared to the only sides roughness but suffers from higher phase variations (Fig. 2(d)) because rapid lateral variations can introduce additional spatial frequencies that interact with the guided mode, potentially leading to stronger or more broadband scattering.

Fig. 3 compares two simulation methods, 3D-FDTD and VarFDTD, for predicting the loss and phase response of a 10 µm long silicon waveguide and extends the analysis to a silicon nitride waveguide under sidewall roughness conditions. Because of the differences in how each method accounts for out-of-plane field components, there are slight discrepancies in the calculated loss (Fig. 3(a)) and phase values (Fig. 3(b)), particularly at higher σ_{RMS}. Both methods, however, show the same qualitative trends: increasing the σ_{RMS} of the sidewall roughness results in larger scattering losses. L_C for side roughness modulates how these scattering centers are distributed along the waveguide axis. In Fig. 3(c) and (d), results are for Silicon Nitride (Si_3N_4) as the core material for the waveguide, which is operated at 780nm wavelength. The figure shows that for comparable σ_{RMS}, Si_3N_4 waveguides can exhibit lower losses as compared to Si-waveguides when the sidewalls (the XZ plane) are rough, particularly if the correlation length is comparable to the waveguide physical dimensions (thickness). Phase shifts also remain lower for Si_3N_4 waveguides as compared to Si counterpart across the whole range of σ_{RMS}.

Fig. 4 presents the loss and phase variation observed in 100 samples of a 10 µm long silicon waveguide, all containing the same sidewall σ_{RMS} but with distinct random placements of the roughness peaks. Although the σ_{RMS} remains constant in all cases, subtle differences in how the peaks and valleys align along the waveguide sidewalls give rise to unique scattering centers that can direct light into radiation modes or shift the effective waveguide index differently. This variability leads to a spread in both the measured loss and phase values across the one hundred samples, demonstrating that nominally identical fabrication processes can yield variations in device performance when random sidewall irregularities are present.

The roughness with a constant value of L_C for a waveguide may not represent the realistic picture as for a longer waveguide, L_C can vary along the length considering an average value of σ_{RMS}. Thus, to average out the randomness of the smaller waveguides, we combined the different random smaller waveguides to make a larger waveguide with the same L_C and σ_{RMS}. Fig. 5 shows the loss and phase variation across 10k distinct 1cm long silicon waveguides, each constructed by randomly stitching together 1000 segments of 10 µm waveguides out of 100 sample geometries with identical σ_{RMS} but different random placements of roughness. Because each 10 µm segment originates from an etch process that might produce unique sidewall features along the waveguide edges, adding a small ensemble of such segments in a random sequence creates an array of possible waveguide realisations. This effect is especially relevant when waveguides become as long as 1cm, which amplifies minor scattering differences that might be negligible in shorter sections for a few random samples. The average loss/phase variation for L_C=50nm is less/more than for waveguides with L_C=100nm across the σ_{RMS} range.

These findings highlight the importance of not only managing the σ_{RMS} and L_C in waveguide fabrication but also recognizing that seemingly identical processes can yield a distribution of device performances once one considers many independent sections. The random variations reflect the reality that large-scale integrated photonics inevitably involves thousands or millions of waveguide elements whose sidewall profiles are not perfectly replicated. Understanding how the concatenation of random segments impacts total loss and global phase is essential when designing longer photonic paths, such as those required for delay lines, sensors, and complex interferometric circuits. Improvements in lithographic resolution, etch chemistry control, and post-processing techniques that smooth sidewalls can reduce the spread of results observed here by driving the σ_{RMS} to sub-nm levels or by suppressing high spatial-frequency irregularities. Nevertheless, even advanced processes cannot entirely eliminate random variations, so computationally efficient methods like FDTD, combined with statistical analysis of device ensembles, provide valuable insights into the expected performance range in real-world photonic systems.

III. Conclusions

Our work shows that the presence of sidewall surface roughness at waveguides can have a significant impact and lead to signal loss and transmission. The results of our large-scale numerical investigations confirm that sidewall roughness remains a key determinant of loss and phase behavior in silicon photonic waveguides. Even when waveguides share identical σ_{RMS} and L_C, randomizing the placement of roughness peaks in each segment introduces notable variability in cumulative scattering. By stitching together multiple 10 µm sections with distinct roughness profiles, we have replicated the construction of practical long waveguides in real photonic circuits. The observed spread in performance emphasizes how variations in fabrication can

lead to inconsistent device characteristics, even under nominally similar process conditions. As silicon photonics continues to expand into new applications, from quantum photonics to advanced computing architectures, a deep understanding of roughness-induced loss and phase variation will be instrumental in pushing the limits of integration density and device functionality.

REFERENCES

[1] Thraskias, C.A., Lallas, E.N., Neumann, N., Schares, L., Offrein, B.J., Henker, R., Plettemeier, D., Ellinger, F., Leuthold, J. and Tomkos, I., 2018. Survey of photonic and plasmonic interconnect technologies for intra-datacenter and high-performance computing communications. *IEEE Communications Surveys & Tutorials*, *20*(4), pp.2758-2783.

[2] Xiang, C., Bowers, S.M., Bjorlin, A., Blum, R. and Bowers, J.E., 2021. Perspective on the future of silicon photonics and electronics. *Applied Physics Letters*, *118*(22).

[3] Chiles, J. and Fathpour, S., 2017. Silicon photonics beyond silicon-on-insulator. *Journal of Optics*, *19*(5), p.053001.

[4] Bogaerts, W., Baets, R., Dumon, P., Wiaux, V., Beckx, S., Taillaert, D., Luyssaert, B., Campenhout, J.V., Bienstman, P. and Thourhout, D.V., 2005. Nanophotonic waveguides in silicon-on-insulator fabricated with CMOS technology. *Journal of Lightwave Technology*, *23*(1), p.401.

[5] Barwicz, T. and Haus, H.A., 2005. Three-dimensional analysis of scattering losses due to sidewall roughness in microphotonic waveguides. *Journal of Lightwave Technology*, *23*(9), p.2719.

[6] Grillot, F., Vivien, L., Laval, S., Pascal, D. and Cassan, E., 2004. Size influence on the propagation loss induced by sidewall roughness in ultrasmall SOI waveguides. *IEEE Photonics Technology Letters*, *16*(7), pp.1661-1663.

[7] Melati, D., Melloni, A. and Morichetti, F., 2014. Real photonic waveguides: guiding light through imperfections. *Advances in Optics and Photonics*, *6*(2), pp.156-224.

[8] Feilchenfeld, N.B., Nummy, K., Barwicz, T., Gill, D., Kiewra, E., Leidy, R., Orcutt, J.S., Rosenberg, J., Stricker, A.D., Whiting, C. and Ayala, J., 2017, March. Silicon photonics and challenges for fabrication. In *Advanced Etch Technology for Nanopatterning VI* (Vol. 10149, p. 101490D). SPIE.

[9] Lee, K.K., Lim, D.R., Luan, H.C., Agarwal, A., Foresi, J. and Kimerling, L.C., 2000. Effect of size and roughness on light transmission in a Si/SiO 2 waveguide: Experiments and model. *Applied Physics Letters*, *77*(11), pp.1617-1619.

[10] Tong, X.C., 2013. Fundamentals and design guides for optical waveguides. *Advanced Materials for Integrated Optical Waveguides*, pp.1-51.

[11] Ladouceur, F., 1997. Roughness, inhomogeneity, and integrated optics. *Journal of lightwave technology*, *15*(6), pp.1020-1025.

[12] Yakuhina, A., Kadochkin, A., Svetukhin, V., Gorelov, D., Generalov, S. and Amelichev, V., 2020, November. Investigation of side wall roughness effect on optical losses in a multimode Si3N4 waveguide formed on a quartz substrate. In *Photonics* (Vol. 7, No. 4, p. 104). MDPI.

[13] Grillot, F., Vivien, L., Laval, S., Pascal, D. and Cassan, E., 2004. Size influence on the propagation loss induced by sidewall roughness in ultrasmall SOI waveguides. *IEEE Photonics Technology Letters*, *16*(7), pp.1661-1663.

[14] Yap, K.P., Delâge, A., Lapointe, J., Lamontagne, B., Schmid, J.H., Waldron, P., Syrett, B.A. and Janz, S., 2009. Correlation of scattering loss, sidewall roughness and waveguide width in silicon-on-insulator (SOI) ridge waveguides. *Journal of Lightwave Technology*, *27*(18), pp.3999-4008.

[15] Morichetti, F., Canciamilla, A., Ferrari, C., Torregiani, M., Melloni, A. and Martinelli, M., 2010. Roughness induced backscattering in optical silicon waveguides. *Physical review letters*, *104*(3), p.033902.

[16] Roberts, S., et al., 2022. Measurements and Modeling of Atomic-Scale Sidewall Roughness and Losses in Integrated Photonic Devices. *Advanced Optical Materials*, *10*(18), p.2102073.

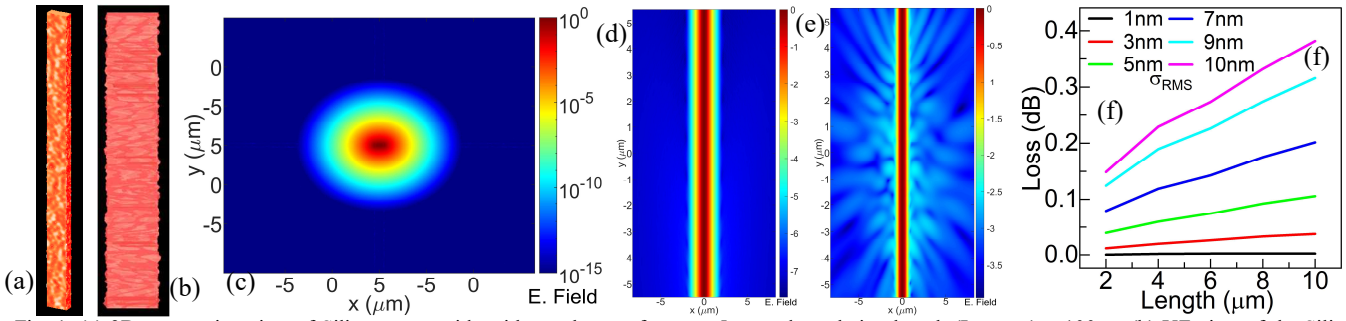

Fig. 1. (a) 3D perspective view of Silicon waveguide with roughness of σ_{RMS} = 5nm and corelation length ($L_{C\ (X/Y/Z)}$) = 100nm (b) YZ view of the Silicon waveguide (c) Mode profile in XY plane (d) Electric field profile of the plane Silicon waveguide without roughness (XZ plane) (e) Electric field profile of the Silicon waveguide with roughness of σ_{RMS} = 1nm and $L_{C\ (X/Y/Z)}$ = 100nm (XZ plane) (f) Loss variation across the length of Silicon waveguide for different σ_{RMS}

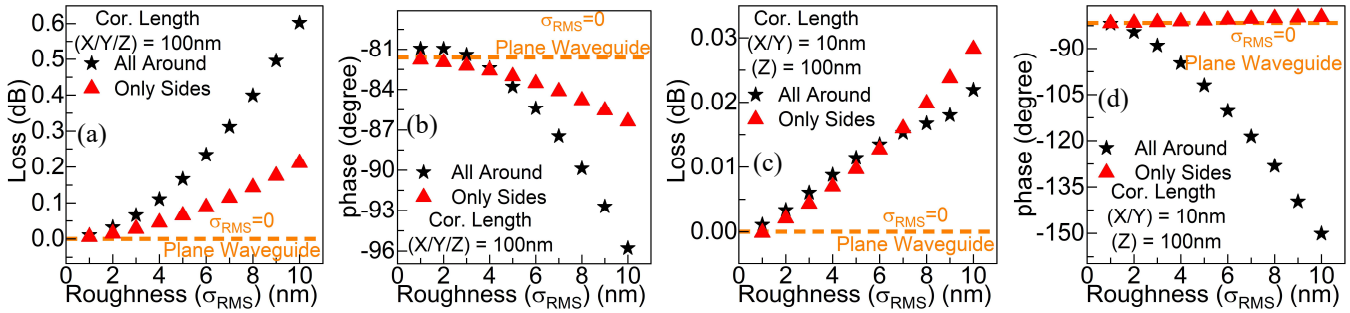

Fig. 2. Loss and Phase variation of a 10μm long Silicon waveguide for different roughness all around and only sides (XZ plane) with various σ_{RMS} values and (a) & (b) $L_{C\ (X/Y/Z)}$ = 100nm; (c) & (d) $L_{C\ (X/Y)}$ = 10nm and $L_{C\ (Z)}$ = 100nm respectively (Full 3D-FDTD)

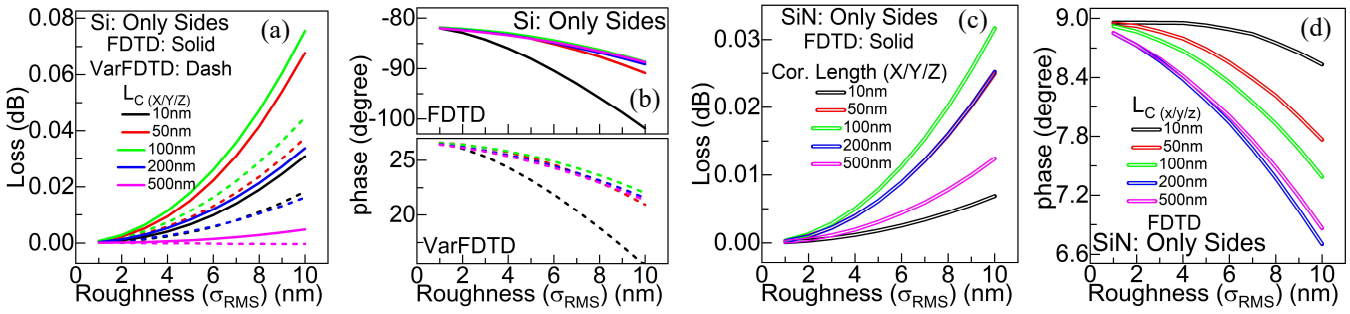

Fig. 3. (a) & (b) Comparing the effect of 3D-FDTD and VarFDTD simulations on loss and phase variation of a 10μm long Silicon waveguide (c) & (d) loss and phase variation of a 10μm long Silicon Nitride waveguide for different roughness (only sides- XZ plane) with various σ_{RMS} and $L_{C\ (X/Y/Z)}$ values, respectively

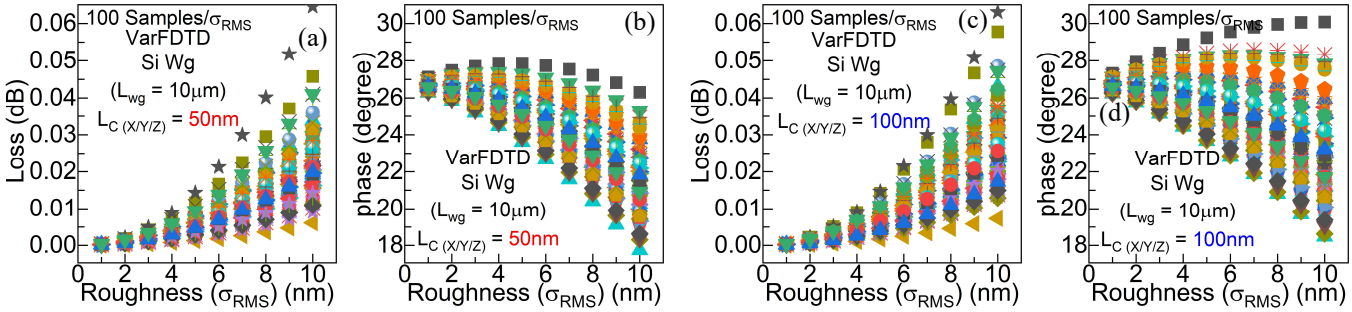

Fig. 4. Loss and phase variation in 100 samples of a 10μm long Silicon waveguide with same roughness (σ_{RMS}) only on the sides (XZ plane) but different randomness (location of roughness peaks changes) (a) & (b) $L_{C\ (X/Y/Z)}$ =50nm and (c) & (d) $L_{C\ (X/Y/Z)}$ =100nm, respectively

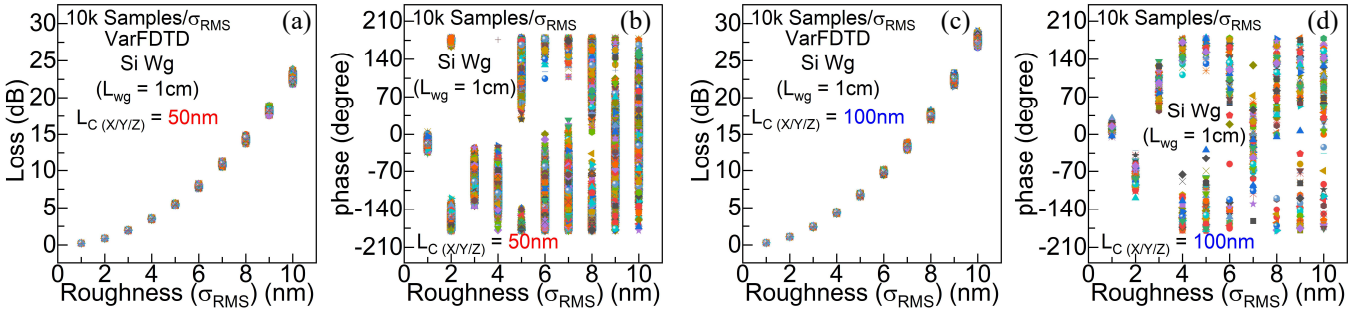

Fig. 5. Loss and phase variation in 10k samples of a 1cm long Silicon waveguide made up of randomly selected (out of 100 samples) 10μm long Silicon waveguide with same roughness (only on the sides (XZ plane)) but different randomness (location of roughness peaks changes) (a) & (b) $L_{C\ (X/Y/Z)}$ =50nm and (c) & (d) $L_{C\ (X/Y/Z)}$ =100nm, respectively [VarFDTD]