Université d'Aix-Marseille

Thèse

pour obtenir le grade de
Docteur de l'université d'Aix-Marseille

École Doctorale :
Sciences pour l'Ingénieur – Mécanique, Physique, Micro et Nanoélectronique
Spécialité : Micro et Nanoélectronique

présentée par
Victor Malherbe

Multi-Scale Modeling of Radiation Effects for Emerging Space Electronics: from Transistors to Chips in Orbit

Modélisation multi-échelle des effets radiatifs pour l'électronique spatiale émergente : des transistors aux puces en orbite

soutenue le 17 décembre 2018 devant le jury composé de :

Pr.	Jean-Luc Autran	IM2NP, Université d'Aix-Marseille	Directeur de thèse
Dr.	Daniela Munteanu	IM2NP, Université d'Aix-Marseille	Co-directrice de thèse
Dr.	Gilles Gasiot	STMicroelectronics	Encadrant industriel
Pr.	Frédéric Wrobel	IES, Université de Montpellier	Président, rapporteur
Pr.	Michael Alles	ISDE, Vanderbilt University	Rapporteur
Dr.	Markus Brugger	EN-EA, CERN	Examinateur

Thèse préparée dans le cadre d'une convention CIFRE entre la société STMicroelectronics et le laboratoire IM2NP (UMR CNRS 7334)

Abstract

The effects of cosmic radiation on electronics have been studied since the early days of space exploration, given the severe reliability constraints arising from harsh space environments. However, recent evolutions in the space industry landscape are changing radiation effects practices and methodologies, with mainstream technologies becoming increasingly attractive for radiation-hardened integrated circuits. Due to their high operating frequencies, new transistor architectures, and short rad-hard development times, chips manufactured in latest CMOS processes pose a variety of challenges, both from an experimental standpoint and for modeling perspectives. This work thus focuses on simulating single-event upsets and transients in advanced FD-SOI and bulk silicon processes.

The soft-error response of 28 nm FD-SOI transistors is first investigated through TCAD simulations, allowing to develop two innovative models for radiation-induced currents in FD-SOI. One of them is mainly behavioral, while the other captures complex phenomena, such as parasitic bipolar amplification and circuit feedback effects, from first semiconductor principles and in agreement with detailed TCAD simulations.

These compact models are then interfaced to a complete Monte Carlo Soft-Error Rate (SER) simulation platform, leading to extensive validation against experimental data collected on several test vehicles under accelerated particle beams. Finally, predictive simulation studies are presented on bit-cells, sequential and combinational logic gates in 28 nm FD-SOI and 65 nm bulk Si, providing insights into the mechanisms that contribute to the SER of modern integrated circuits in orbit.

Keywords—cosmic ray, single-event upset (SEU), single-event transient (SET), soft-error rate (SER), compact model, Monte Carlo simulation, fully-depleted silicon on insulator (FD-SOI), 28 nm, bipolar amplification, floating body effect.

Résumé

En raison de leur impact sur la fiabilité des systèmes, les effets du rayonnement cosmique sur l'électronique ont été étudiés dès le début de l'exploration spatiale. Néanmoins, de récentes évolutions industrielles bouleversent les pratiques dans le domaine, les technologies standard devenant de plus en plus attrayantes pour réaliser des circuits durcis aux radiations. Du fait de leurs fréquences élevées, des nouvelles architectures de transistor et des temps de durcissement réduits, les puces fabriquées suivant les derniers procédés CMOS posent de nombreux défis. Ce travail s'attelle donc à la simulation des aléas logiques permanents (SEU) et transitoires (SET), en technologies FD-SOI et bulk Si avancées.

La réponse radiative des transistors FD-SOI 28 nm est tout d'abord étudiée par le biais de simulations TCAD, amenant au développement de deux modèles innovants pour décrire les courants induits par particules ionisantes en FD-SOI. Le premier est principalement comportemental, tandis que le second capture des phénomènes complexes tels que l'amplification bipolaire parasite et la rétroaction du circuit, à partir des premiers principes de semi-conducteurs et en accord avec les simulations TCAD poussées.

Ces modèles compacts sont alors couplés à une plateforme de simulation Monte-Carlo du taux d'erreurs radiatives (SER), conduisant à une large validation sur des données expérimentales recueillies sous faisceau de particules. Enfin, des études par simulation prédictive sont présentées sur des cellules mémoire et portes logiques en FD-SOI 28 nm et bulk Si 65 nm, permettant d'approfondir la compréhension des mécanismes contribuant au SER en orbite des circuits intégrés modernes.

Mots-clés—rayon cosmique, aléas logiques (SEU, SET, SER), modèle compact, simulation Monte-Carlo, silicium sur isolant totalement déserté (FD-SOI), 28 nm, amplification bipolaire, effet de corps flottant.

À la mémoire de mon père.

Acknowledgments

Regardless of the audience this thesis may or may not reach, simply being given the opportunity to perform research, and complete a body of work that I can be proud of, has been an extraordinary experience in its own right. Within the framework of this industrial "CIFRE" PhD, I seem to have found a sweet spot between engineering and science. This odd field of radiation effects, it appears, is quite worth the little detour.

Prime contributors are of course immediate colleagues. First and foremost, I would like to warmly thank my PhD advisors for their guidance throughout the course of this PhD program. To Jean-Luc my supervisor at IM2NP lab, thank you for all the fruitful interactions surrounding my work, for the much-needed help with Geant4 nuclear simulations, and in short for the brilliant partnership with the team at ST. To Gilles my supervisor at ST, thank you for having tutored my physical intuition from my very first steps in the field—and coping so remarkably well with my *occasional* tendency to procrastinate. And he might not have been one of my official advisors, but I would also like to thank Philippe R. for all the pep talks and for taking on the role of informal coach every now and then. Steering a boat with such a motley crew is certainly not the easiest of tasks.

For the wealth of silicon data that they spoil us with, I am grateful to the experimentalists in and outside the team. Hats off to Dimitri, Serge, Vincent, plus collaborators at EASii, IROC, HIREX, IM2NP, working on beam tests both day and night. I would also like to extend my thanks to the circuit designers, Fady, Sylvain, Martin, Maximilien, Guénolé a.k.a. my weekend companion in misfortune, Capucine, Olivia, Anna, Calogero, Dominique from Maya, for their unwavering dedication to the cause—that is of course, sleep deprivation before tape-out to provide us with actual chips to irradiate. Many thanks also to my other co-workers Benjamin, Cyril, Valérie, and Jean-Marc my *reluctant at first but actually not so dreadful* cubicle-mate. And finally I am greatly indebted to the "software people" who have helped grow TIARA, the soft-error simulator described in this manuscript, into the industrial tool that it is today: first Philippe C. from Immotronic and, to an immense extent, Thomas—without whom TIARA would be but an R&D curiosity, unusable of anyone except the one PhD student behind its obscure code. "Physicist meets software specialist", as it turns out, does not always end in mutual incomprehension.

Another kind of acknowledgments is those I wish to direct to the institutions that played a role in forging my scientific mind and critical thinking over the years. I received very broad education in nanotechnology during my time at University College London; it may have taken me an entire MSc thesis to understand how *stiff* semiconductor equations are, but it proved essential to my later understanding of charge transport in silicon. I also want to recognize my engineering school Supélec Paris and the quality

training in electrical engineering that I was given—and beyond that a remarkable course in philosophy of science, the teachings of which I am still struggling to truly absorb and live by. Also as far as memorable lessons go, being taught Maxwell's equations on my eighteenth birthday was something that I will hardly ever recover from; throughout my *maths sup – maths spé* training at Chateaubriand Rennes, our relentless, nay frantic teachers played an honorable part, and I mean here to formally acknowledge their crucial role in laying the foundation of my scientific knowledge. In retrospect, the rest of my education has merely been to stack up more modern theories onto that solid basis.

Now, there are those people whose direct influence cannot be measured, but whose essential role I could not fail to mention. Anyone who has had to wrestle against the ups and downs of a long-term undertaking knows about the importance of friends and their unconditional support. I shall refrain from the extensive task of listing their names here—they will know. Long story short, I am simply grateful to them for all the travel, the music, and for all the rock climbing and other moutain-related singularities in my life. And of course, for the occasional chitchat and beverage and the work–life balance that I seem to have achieved, thanks to everyone from Rennes, Paris, London, Grenoble and elsewhere.

Last but not least, my final acknowledgments are to my family. My mother and sister may not crave for quantum mechanics textbooks, but in ways that I shall probably be the last to decipher, they have contributed their very own intellectual pieces to the puzzle. And by nurturing my scientific curiosity from the youngest age, my father has beyond doubt played the greatest role in crafting the physics-loving engineer that I have come to be. For this, and for so many other things that I carry with me always, I am ever grateful.

Résumé détaillé

Contexte : évènements singuliers et électronique spatiale

Contrairement à l'environnement terrestre qui, grâce au champ magnétique et à l'atmosphère de la Terre, est à l'abri de la plupart des radiations en provenance de l'espace, l'environnement spatial contient de nombreuses particules énergétiques susceptibles d'affecter l'électronique. Les effets de ce rayonnement cosmique sur l'électronique spatiale se scindent généralement en deux larges catégories que sont les effets singuliers (*Single-Event Effects* ou SEE) et les effets de dose, qu'elle soit ionisante ou non. Ce travail se focalise sur les SEEs non destructifs, ou *soft errors*, plus particulièrement les aléas logiques permanents (*Single-Event Upset*, ou SEU) et transitoires (*Single-Event Transient*, ou SET). Ces phénomènes sont modélisés par des approches multi-échelle qui permettent de rendre compte du dépôt de charge initial par les particules énergétiques, de la physique du semiconducteur qui gouverne la réponse des composants élémentaires, ainsi que des effets au niveau circuit, dans le domaine analogique ou numérique. La première brique est étudiée à l'aide d'outils Monte-Carlo nucléaires tels que Geant4 [1] ou MCNP [2] ; la réponse du silicium fait généralement l'objet de simulations *Technology Computer-Aided Design* (TCAD) ; enfin, dans le domaine électrique les simulateurs circuit basé sur SPICE sont privilégiés pour les études au niveau porte logique.

À la croisée de ces différentes couches d'abstraction, dans les années 2000 des outils de calcul du taux d'erreurs radiatives (*Soft-Error Rate*, ou SER) par méthode de Monte-Carlo ont vu le jour, tels que MRED [3] ou MC-ORACLE [4]. Ces codes de simulation déclinent de nombreux choix de modélisation mais ont pour caractéristique commune d'apporter un traitement simplifié de la brique "semiconducteur" : la résolution directe des équations de transport de charge par des approches de type TCAD étant très coûteuse en temps de calcul, elle ne se prête pas à un calcul par "force brute" du SER qui nécessite d'intégrer la réponse des dispositifs sur des milliers, voire des millions d'impacts de particules. Ainsi, le simulateur *Tool suIte for rAdiation Reliability Assessment* (TIARA), développé par STMicroelectronics et l'IM2NP et dont il est question dans ce travail, procède aux calculs de physique du semiconducteur grâce à des modèles compacts dédiés qui seront examinés par la suite.

Le contexte industriel de ce travail de thèse s'articule autour des évolutions récentes dans les deux domaines que sont le spatial et le marché du semiconducteur. Tout d'abord, nous assistons à un essor des applications spatiales "de masse" en orbite terrestre, parmi lesquelles figurent par exemple des projets de constellations de milliers de satellites visant à fournir une couverture Internet mondiale ; de tels volumes inédits de satellites mettent les problématiques de radiation sur le devant de la scène – et en particulier la modélisation comme outil de prédiction plus flexible que les mesures expérimentales

sous faisceau – tandis que dans le même temps les objectifs de performance (et donc de rentabilité) ne peuvent être atteints que par des compromis agressifs vis-à-vis de la fiabilité et de la robustesse des circuits. Ceci conduit l'industrie spatiale à souscrire de plus en plus aux technologies commerciales pour réaliser des circuits durcis, au lieu de technologies certes éprouvées mais vieillissantes. Le deuxième axe industriel de cette thèse s'inscrit donc naturellement dans le contexte de miniaturisation des circuits intégrés, et notamment l'adoption des nouvelles architectures de transistors que sont les technologies *Fin Field-Effect Transistor* (FinFET) et silicium sur isolant totalement déserté (*Fully-Depleted Silicon On Insulator*, ou FD-SOI), dont les morphologies très éloignées du transistor planaire en silicium massif (*bulk Si*) conduisent à des propriétés nouvelles vis-à-vis des *soft errors*.

Les objectifs de ce travail consistent donc à accompagner le développement des plateformes spatiales récentes chez STMicroelectronics, à savoir l'offre *bulk Si* en technologie 65 nm, et le FD-SOI au nœud 28 nm. Pour ce faire, les spécificités du FD-SOI doivent être étudiées par le biais de simulations TCAD afin de comprendre les mécanismes de collection de charge sous-jacents. Cela doit conduire au développement de modèles compacts dédiés pour permettre des prédictions précises du SER de circuits en technologie FD-SOI dans l'outil Monte-Carlo TIARA. Un aspect plus technique va également consister en l'adaptation, la maintenance, et l'amélioration du code TIARA, en particulier pour permettre la simulation du FD-SOI dans un outil initialement prévu pour opérer sur les technologies *bulk*. Plus généralement, les développements logiciels viseront à automatiser au maximum le flot de simulation et son intégration dans les outils CAD standard. Enfin, les tâches de recherche et développement sus-citées ont pour but de répondre au besoin initial formulé, qui est de contribuer à la réalisation de circuits durcis pour des applications spatiales modernes. À cet égard, la modélisation prédictive doit permettre d'explorer plus de scénarios que ce qui est permis expérimentalement – particulièrement pour la technologie FD-SOI 28 nm qui est nativement très robuste, donc difficile à caractériser sous faisceau. Enfin, tout autant que la prédiction du SER en orbite de circuits déjà conçus et figés, les outils de modélisation devront être mis en œuvre pour édicter des règles de conception robuste pour se prémunir des phénomènes SEU et SET afin de consolider les plateformes spatiales en 65 et 28 nm.

TIARA : une plateforme de simulation Monte-Carlo du SER

Le code de simulation TIARA, auquel ces travaux contribuent, a été initié à ST et l'IM2NP à la fin des années 2000 [5]. Le programme opère un simulateur SPICE qui réalise des injections de courant parasite dans des cellules logiques CMOS pour prédire la réponse à un impact de particule ; chaque simulation électrique s'exécute en quelques secondes, ce qui permet d'aboutir par méthode de Monte-Carlo (intégration probabiliste) à un taux d'erreur ou une section efficace radiative en quelques minutes ou heures. Les processus physiques en amont de l'injection de courant, à savoir le dépôt des charges et leur collection, font l'objet de différents modules[1] décrits sur la Figure 1.

Tout d'abord, étant donnés un schéma électrique et son dessin physique (une *netlist* et un *layout*), le simulateur crée une structure de simulation qui représente le circuit

[1]On entend ici par "module" les principales entités qui structurent TIARA, qui correspondent à autant de classes – au sens programmation orientée objet – de son code C++. À l'heure actuelle, TIARA comptabilise environ 53 000 lignes de code C++, ainsi que de nombreux scripts shell et Python pour faciliter la mise en place des simulation et les analyses ultérieures.

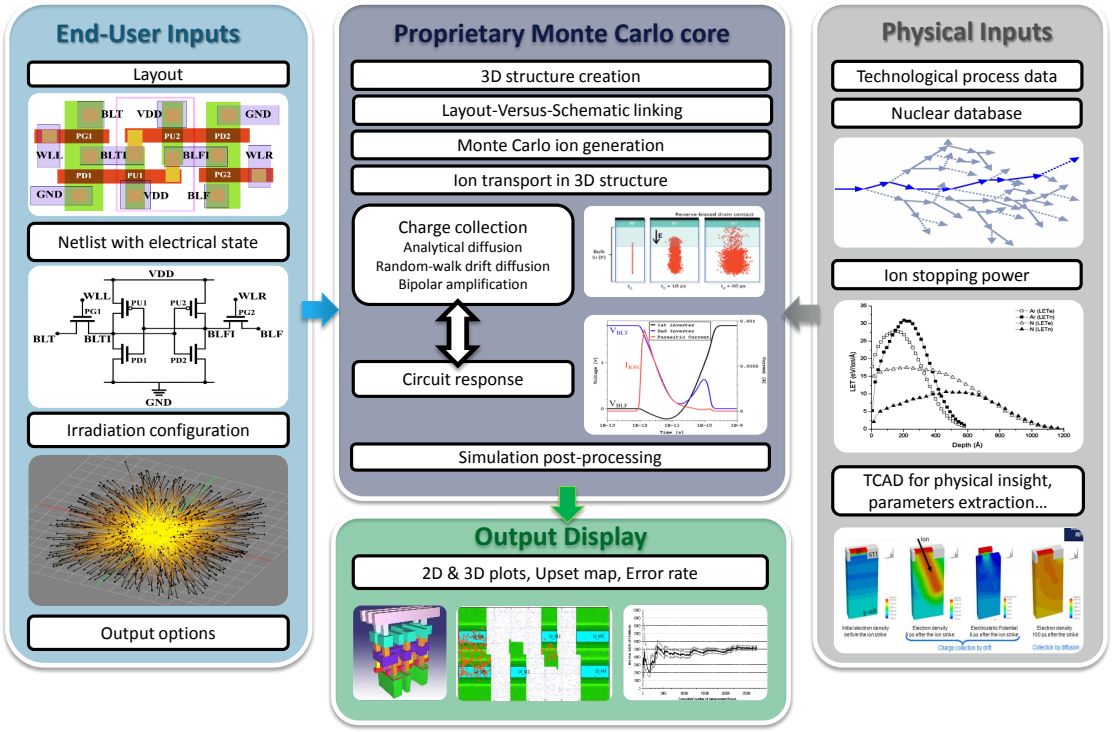

Figure 1: Schéma d'architecture des principaux composants et entrées–sorties de TIARA.

dans les domaines physique et électrique : un module "`Builder3D`" se charge de lire les polygones (bidimensionnels) du *layout* (au format GDS) et combine les différents calques CAD du GDS par des opérations booléennes. Cela permet au `Builder3D` de créer toute la géométrie 2D associée au circuit, qui est ensuite extrudée sur la dimension verticale moyennant des informations sur l'empilement du procédé de manufacture. On obtient ainsi un modèle 3D complet du circuit intégré étudié. La *netlist* du circuit est alors lue (au format CDL avec ou sans parasites au format DSPF), et s'ensuit une étape cruciale opérée par un module nommé "`LVSmatcher`" : les matrices de connectivité de la *netlist* et du *layout* sont extraites, puis mises en correspondance, à la manière d'une opération de *Layout Versus Schematic*, ou LVS. Ceci est nécessaire pour par la suite être capable d'injecter des courants dans une simulation électrique, en reliant automatiquement les coordonnées physiques d'impacts de particules à des nœuds du schéma électrique. Ces deux modules de `Builder3D` et `LVSmatcher`, absents de TIARA dans les travaux antérieurs à cette thèse, constituent un ajout significatif fourni par l'auteur et des collègues informaticiens.

Passée la création de la représentation interne du circuit, TIARA rentre dans la boucle Monte-Carlo proprement dite : le module "`Irradiator`" tire des ions aléatoirement suivant l'environnement radiatif configuré, en d'autres termes, il échantillonne les variables aléatoires que sont le type et l'énergie des noyaux, ainsi que leur point de départ et leur direction. Les scénarios d'irradiation couverts sont les environnements terrestres – particules alpha isotropes ou neutrons atmosphériques verticaux, dont les reculs sur silicium sont lus dans des bases de données Geant4 dédiées – ainsi que les environnements spatiaux synthétiques ou réels : ions lourds unidirectionnels et mono-énergétiques pour simuler une manipulation sous faisceau, ions isotropes à spectre large en orbite, et produits de réactions sous protons, encore une fois gérés par des bases de données Geant4. Dès lors, le "`Raytracer`" est en mesure de transporter ces noyaux

dans la structure 3D du circuit, pour obtenir les charges déposées le long des trajectoires d'ions. Le `Raytracer` opère dans la *Continuous Slowing Down Approximation* (CSDA), propageant les ions en ligne droite avec une atténuation d'énergie issue du *Linear energy Transfer* (LET) de SRIM [6] pour tous les couples ion–matériau nécessaires.

En se basant sur les traces d'ions en sortie du `Raytracer`, le module "`Collector`" simule alors le transport des porteurs de charge dans les semiconducteurs, ce qui donne lieu à des courants parasites à injecter dans une simulation SPICE. En procédé *bulk*, TIARA implémente le modèle de diffusion–collection provenant de [7], qui s'appuie sur une formule analytique de diffusion ambipolaire des porteurs dans les zones neutres de substrat et de caissons. Cette hypothèse n'est plus valide en FD-SOI, où le canal totalement déserté implique que le transport des porteurs est désormais dominé par le courant de dérive dans le champ électrique. Indépendamment du modèle de collection, les APIs de TIARA permettent alors aux courants issus du `Collector` de s'insérer dans une simulation électrique, grâce à un module "`CircuitSolver`" qui opère le simulateur Eldo [8] par le biais d'appels système. La simulation électrique peut être effectuée d'une seule traite une fois terminés les calculs du `Collector`, ou au contraire pas à pas (grâce au mode interactif d'Eldo), ce qui permet de tenir compte de la rétroaction du circuit sur la collection de charge. Le mode d'exécution couplé fait ainsi partie des améliorations majeures apportées à TIARA au niveau physique.

Pour finir, les résultats de simulation sont analysés. Les post-traitements s'effectuent à deux niveaux, grâce à un module "`Analyzer`" qui se charge à la fois d'analyses microscopiques et macroscopiques : pour chaque évènement, les formes d'onde simulées sont analysées pour enregistrer d'éventuelles signatures de SEU ou SET, et ces données sont ensuite agrégées afin d'extraire des quantités statistiques telles que la section efficace ou le SER du circuit. Pour finir cette présentation, le module dit "`Driver`" gère le séquençage de toutes les opérations mentionnées, c'est-à-dire l'enchaînement entre les modules, que ce soit pour une simulation en local sur une machine unique, ou distribuée sur une ferme de calcul grâce à l'environnement *Load Sharing Facility*, ou LSF [9]. En terme de comparaison avec l'état de l'art, sur les aspects physiques la principale limitation de `TIARA` réside dans l'approche de type "bases de données" pour traiter les interactions nucléaires – comme c'est le cas de nombreux codes industriels ou commerciaux, contrairement aux codes académiques généralement plus fondamentaux. Les modèles de collection dédiés en *bulk* et FD-SOI, à l'issue de cette thèse, constituent le principal atout de TIARA vis-à-vis de modèles plus anciens tels que les critères de charge critique simplifiés.

Modélisation des SEEs au niveau transistor

Afin d'appréhender les spécificités de la technologie FD-SOI 28 nm sous radiations, des simulations TCAD ont été opérées sur des structures de transistors unitaires [10]. Celles-ci mettent tout d'abord en évidence les volumes sensibles très confinés, grâce à l'isolation électrique fournie par le BOX et la minceur (typiquement 10 nm) du film de silicium actif. Ainsi, la collection de charge s'effectue principalement sous la grille du transistor [11], plus particulièrement à la jonction canal–drain. D'autre part, l'effet d'amplification bipolaire parasite inhérent aux technologies SOI [12] reste modéré en FD-SOI 28 nm : le champ électrique qui règne à l'intérieur du canal expulse rapidement les porteurs majoritaires responsables de l'injection de courant supplémentaire par la source, si bien que pour des LETs typiques (1 $\mathrm{MeV \cdot cm^2/mg}$ et au-delà), le rapport entre la charge collectée et déposée reste inférieur à 3.

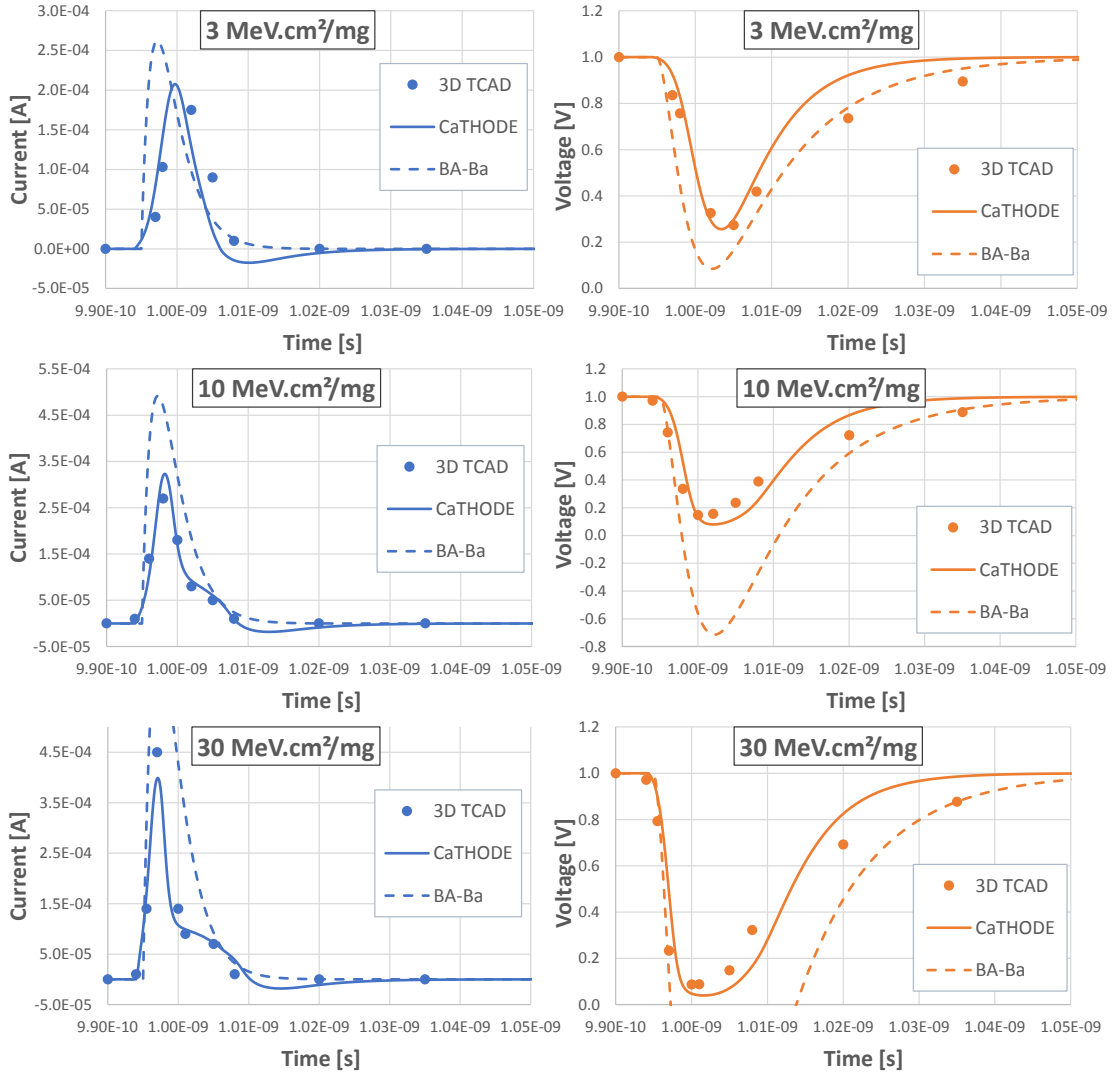

Figure 2: Comparaison entre les modèles de réponse FD-SOI avec des simulations TCAD 3D *mixed-mode* (impacts sur le nMOS d'un inverseur pour plusieurs LETs).

Ces observations ont premièrement été encodées dans un modèle comportemental pour simuler les courants de radiations en FD-SOI dans TIARA [13]. *Bipolar Amplification Basics*, ou BA-Ba, calcule la charge collectée lors d'un impact sur un MOSFET bloqué en multipliant la charge déposée par une amplification bipolaire tabulée en fonction du LET, puis la module par une efficacité de collection fonction de la position d'impact. Ces deux fonctions précalculées s'inspirent des résultats obtenus en TCAD, à savoir un gain bipolaire maximal à bas LET et tendant vers 1 en injection forte, et une extraction maximale à la jonction canal–drain et qui s'atténue en s'éloignant de part et d'autre. Le modèle est ajusté dans le voisinage du point de fonctionnement *Process–Voltage–Temperature* (PVT) nominal, mais par ailleurs, sa principale limitation réside dans le fait qu'il n'est pas dynamiquement couplé au circuit : la charge collectée obtenue est injectée comme une forme d'onde en double exponentielle dont les temps de montée et descente sont fixés à quelques picosecondes (là encore sur la base de résultats TCAD). Dès lors, la rétroaction circuit est ommise, et BA-Ba se cantonne à la simulation SEU : le modèle peut prédire l'occurrence d'une chute de tension, mais puisque le courant radiatif n'est pas modulé par la suite, la durée d'un SET n'est pas correctement capturée.

Le deuxième modèle de collection développé dans le cadre de cette thèse se dénomme *Carrier Transport with Humble One-Dimensional Equations*, ou CaTHODE [14]. L'idée centrale du modèle est de procéder à une résolution directe des équations de semiconducteurs, sur une unique dimension spatiale : le mouvement des charges en FD-SOI est essentiellement latéral, et de ce fait une simulation autoconsistante (électrostatique couplée au transport) le long de l'axe source–canal–drain des MOSFETs impactés doit pouvoir capturer les principaux phénomènes physiques en des temps d'exécution tout à fait raisonnables.

CaTHODE considère donc le modèle de dérive–diffusion, en tenant compte de modèles physiques similaires à ceux disponibles en TCAD : la structure de bande intègre les phénomènes de *BandGap Narrowing* (BGN), essentiels pour capturer correctement le gain bipolaire en présence d'implants fortement dopés ; la recombinaison Schockley–Read–Hall (SRH) se manifeste aux longues échelles de temps, et la recombinaison Auger est un contributeur important à fort LET, particulièrement en cœur de trace d'ion ; la génération par avalanche (ionisation par impact) est susceptible de multiplier la charge collectée sous de forts champs électriques ; enfin le passage de l'ion induit des paires électron–trou dans une certaine zone du dispositif et à un certain instant. Le système d'équations différentielles est discrétisé suivant les procédures couramment employées dans les outils TCAD, notamment le schéma de Scharfetter–Gummel pour les équations de courant. Le système non-linéaire qui en résulte est alors inversé par méthode de Newton, grâce à une factorisation *Lower–Upper* (LU). La bibliothèque d'algèbre linéaire Eigen [15] est utilisée à cette fin. Lors d'un impact ionisant, TIARA génère donc un maillage 1D pour chaque transistor impacté, auquel il assigne un profil de dopage, et une simulation transitoire peut débuter, produisant un courant de radiation injecté dans Eldo. Les conditions aux limites sur ces maillages sont mises à jour dynamiquement pour tenir compte de la rétroaction du circuit, et ainsi CaTHODE permet de simuler finement les phénomènes SET. Ceci est mis en évidence sur la Figure 2, où le nMOS d'un inverseur est irradié à différents LETs : à fort LET, le modèle comportemental BA-Ba surestime l'extraction de charge – et donc l'amplitude et la durée du SET – car le courant n'est pas modulé par la chute de tension aux bornes du composant. CaTHODE tient nativement compte de cet effet, car le champ électrique de la simulation 1D évolue avec la réponse du circuit.

En résumé, à l'issue de ce travail, TIARA a donc été muni de deux nouveaux modèles de collection en FD-SOI, disposant chacun d'un domaine de validité qui lui est propre. L'un comme l'autre respectent parfaitement les critères de compacité pour la simulation SER Monte-Carlo, BA-Ba s'exécutant quasi-instantanément et CaTHODE tournant typiquement en quelques secondes par impact.

Simulation des SEEs au niveau cellule logique

La chaîne complète de simulation SER bâtie en FD-SOI a tout d'abord été validée sur des données expérimentales. La calibration des modèles compacts de collection s'effectue en confrontant les simulations à des résultats recueillis sous faisceau sur des véhicules de test embarquant des mémoires SRAM en 28 nm. Ainsi, les paramètres libres de BA-Ba et CaTHODE sont ajustés pour reproduire au mieux les courbes de section efficace sur un panel de cellules et de rayonnements. Suite à cette procédure de calibration, TIARA peut être utilisé pour examiner les zones sensibles sur des cellules mémoire, comme le montre la Figure 3 où les zones sensibles d'une SRAM à bas LET résultent

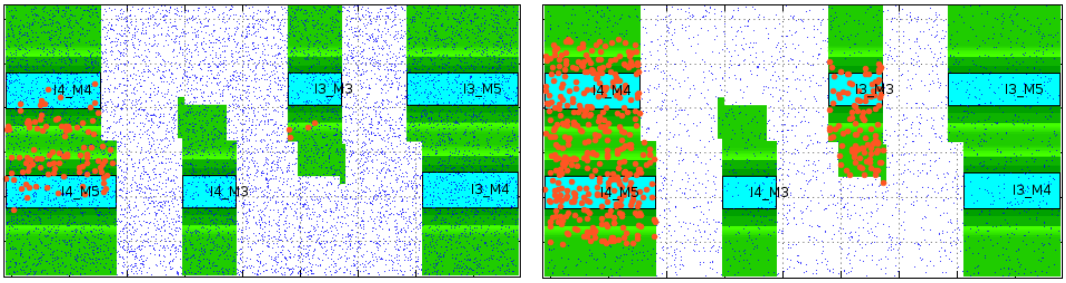

(a) (b)

Figure 3: Cartographie d'erreurs simulée sur une cellule SRAM, à 1.83 MeV·cm^2/mg (a) et 60 MeV·cm^2/mg (b). Vert: silicium actif; bleu (resp. rouge): impacts sans SEU (resp. avec).

d'une compétition entre l'épaisseur de silicium (due à la morphologie des *spacers* et des implants épitaxiés) induisant un dépôt de charge non uniforme, et l'intensité du champ électrique (déterminé par la position des jonctions) qui influe sur l'extraction de charge. La section efficace de cette cellule mémoire typique est environ cent fois plus faible que la même cellule manufacturée en procédé *bulk*.

Une fois validée la chaîne de simulation sur des cellules mémoire simples, TIARA peut être utilisé pour mener à bien des études sur des cellules plus complexes, lesquelles peuvent être difficiles à tester de manière exhaustive en expérimental. Comme illustré sur la Figure 4, le simulateur est en mesure d'identifier la contribution SEU de chaque transistor d'une bascule D en FD-SOI 28 nm, selon l'état électrique de la cellule et l'énergie du rayonnement incident. De telles données sont difficilement obtensibles sur les structures de test généralement embarquées lors des tests accélérés, et dans cette technologie en particulier la robustesse native implique de longs temps d'exposition et/ou une faible statistique d'erreur. À cet égard, TIARA permet une observabilité accrue, et de nombreuses cellules peuvent être évaluées sous de multiples PVTs.

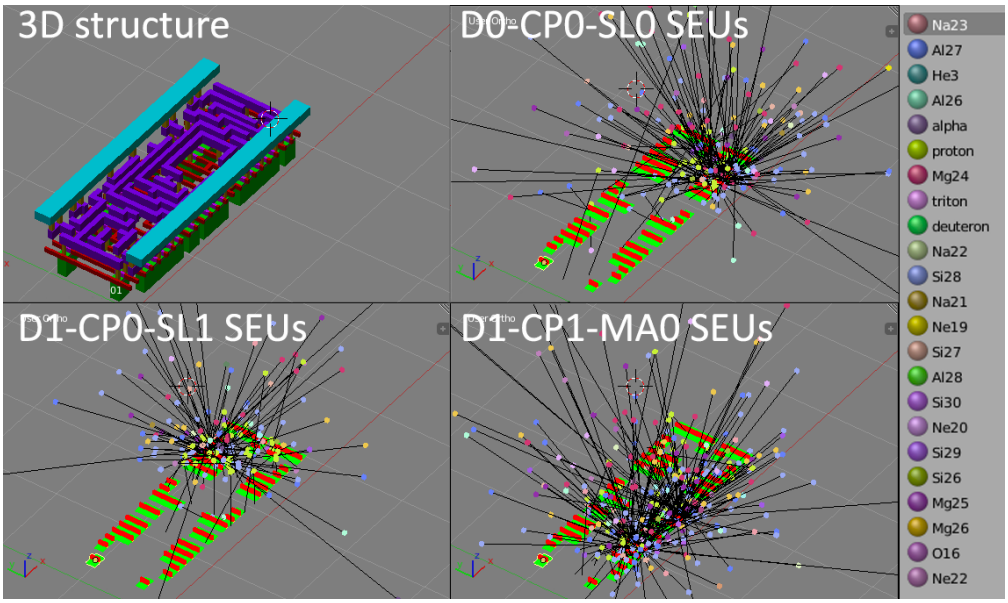

Figure 4: Projections sous protons sur une bascule D en 28 nm : visualisation 3D des reculs induisant un SEU à 100 MeV, pour différents états électriques.

Des études SET ont également été menées sur des portes combinatoires. De larges explorations dans l'offre de cellules standard en FD-SOI 28 nm permettent de distinguer les portes logiques les plus à même de générer des transitoires, généralement du fait du faible courant qui pilote leurs nœuds internes (architecture de la porte) ou leur nœud de sortie (*output drive*). De telles cellules peuvent alors être écartées de l'offre de cellules robustes. En 65 nm, une étude approfondie portant sur la génération et la transmission de SET dans les arbres d'horloge a également été conduite [16]. De tels transitoires ont en effet une faible probabilité de masquage, et peuvent mener à des *Single-Event Functional Interrupts* (SEFI) puisque l'arbre d'horloge pilote de nombreuses bascules du circuit. Là encore, cette étude a débouché sur un ensemble de règles de conception robuste basées sur une sélection pragmatique des cellules de l'offre standard. Enfin, des projections de SER en orbite permettent de quantifier la robustesse des technologies FD-SOI 28 nm et *bulk* 65 nm en environnement réel, pour des satellites à destination d'orbites basses (*Low Earth Orbit*, LEO) ou géostationnaires (*Geostationary Earth Orbit*, GEO). Ainsi, en GEO le FD-SOI 28 nm procure une réduction de SER d'un facteur 50 environ comparé au *bulk* (sur des cellules SRAM en 28 nm). En LEO, le SER du FD-SOI 28 nm est dominé par les protons haute énergie, pour des conditions nominales de blindage et de climat spatial.

Conclusion et perspectives

Une grande part de ce travail de thèse aura donc consisté à développer des modèles compacts pour permettre la simulation SER en technologie FD-SOI : en se basant sur des études TCAD, deux modèles innovants ont été proposés pour calculer les courants induits par particules ionisantes en FD-SOI, sous des temps de calculs compatibles avec la simulation Monte-Carlo nécessitant l'évaluation de plusieurs milliers d'impacts. Le simulateur TIARA a par ailleurs été maintenu et mis à jour pour accueillir ces nouveaux modèles de collection, suite à quoi les modèles de réponse ont été calibrés en confrontant la simulation à des données expérimentales d'irradiation de cellules mémoire 28 nm.

Ces capacités de simulation prédictive ont alors été exploitées pour examiner finement les zones sensibles de cellules mémoire simples, et ont été appliquées à l'étude des phénomènes SEU sur des cellules séquentielles plus complexes telles que des bascules, où la mesure expérimentale ne peut être exhaustive. En outre, les aléas logiques de type SET ont été considérés sur des portes combinatoires en FD-SOI 28 nm et en *bulk* 65 nm, conduisant à édicter des règles de conception robuste pour se prémunir de ces phénomènes transitoires. Enfin, des projections en environnement réel ont permis d'établir le taux d'erreur rendu par de telles cellules en orbite.

Au niveau scientifique, ce travail pourra être étendu en considérant différents points d'amélioration de la chaîne de simulation Monte-Carlo : la physique du dépôt des charges, que ce soit par interaction nucléaire ou ionisation directe, pourra tenir compte de nouveaux matériaux (par exemple le germanium dans le canal en technologies fortement intégrées) et types de rayonnement (notamment les protons basse énergie, de trajectoire non-rectiligne). Les modèles de collection pourront s'atteler à la technologie FinFET dans le futur. D'un point de vue ingénierie, tous ces modèles et outils gagneront à être intégrés toujours plus étroitement au flot de conception, afin de standardiser autant que faire se peut les tâches de durcissement des circuits. Par ailleurs, des ponts restent encore à bâtir avec les systèmes d'injection de fautes (simulateurs numériques), afin de réaliser un passage à l'échelle pour permettre l'analyse systématique de systèmes complexes.

Contents

List of Figures

List of Tables

Acronyms

AABB	Axis-Aligned Bounding Box
AMS	Analog and Mixed Signal
amu	atomic mass unit
API	Application Programming Interface
BA-Ba	Bipolar Amplification Basics
BEOL	Back End Of Line
BGN	BandGap Narrowing
BIST	Built-In Self Test
BJT	Bipolar Junction Transistor
BOX	Buried OXide
BTE	Boltzmann Transport Equation
BVH	Bounding Volume Hierarchy
CAD	Computer-Aided Design
CaTHODE	Carrier Transport with Humble One-Dimensional Equations
CCS	Carrier-Carrier Scattering
CDF	Cumulative Density Function
CDL	Circuit Design Language
CME	Coronal Mass Ejection
CMOS	Complementary Metal Oxide Semiconductor
COTS	Commercial Off-The-Shelf
CRÈME	Cosmic Ray Effects on Micro-Electronics
CSDA	Continuous Slowing Down Approximation
CSV	Comma-Separated Values
DC	Direct Current
DD	Displacement Damage
DFF	D-type Flip-Flop
DFT	Design For Testability
DICE	Dual-Interlocked storage CEll
DoE	Design of Experiments
DRAM	Dynamic Random Access Memory
DRC	Design Rule Check
DRM	Design Rule Manual
DSPF	Detailed Standard Parasitic Format
DUT	Device Under Test
ECC	Error Correction Code

EDA	Electronic Design Automation
ESD	ElectroStatic Discharge
FD-SOI	Fully-Depleted Silicon On Insulator
FEOL	Front End Of Line
FET	Field Effect Transistor
FinFET	Fin Field Effect Transistor
FLUKA	FLUktuierende KAskade
FO	Fan-Out
FPGA	Field-Programmable Gate Array
FWHM	Full Width at Half Maximum
GCR	Galactic Cosmic Ray
GDS	Graphic Database System
Geant4	GEometry ANd Tracking
GEO	Geosynchronous Earth Orbit
GPL	GNU General Public License
GUI	Graphical User Interface
HSSL	High-Speed Serial Link
HZE	High Z and Energy
IC	Integrated Circuit
ICRU	International Commission on Radiation Units and Measurements
IMD	InterMetal Dielectric
IO	Input–Output
IP	Intellectual Property
IRPP	Integral Rectangular ParallelePiped
IRT	Intel Radiation Tool
ISS	International Space Station
JEDEC	Joint Electron Device Engineering Council
LANSCE	Los Alamos Neutron Science CEnter
LEO	Low Earth Orbit
LET	Linear Energy Transfer
LSF	Load Sharing Facility
LU	Lower–Upper
LVS	Layout Versus Schematic
LVT	Low Voltage Threshold
MCNP	Monte Carlo N-Particle
MCU	Multiple-Cell Upset
MHC	Model-Hardware Correlation
MOSFET	Metal Oxide Semiconductor Field Effect Transistor
MRED	Monte Carlo Radiative Energy Deposition
MUSCA SEP3	MUlti-SCAles Single Event Phenomena Predictive Platform
nMOS	n-channel MOSFET
P&R	Place & Route
PB	Poly-Bias
PCB	Printed Circuit Board

PCell	Parametrized Cell
PD-SOI	Partially-Depleted Silicon On Insulator
PDE	Partial Differential Equation
PDF	Probability Density Function
PDI	Proton Direct Ionization
PDK	Process Design Kit
PHITS	Particle and Heavy Ion Transport System
PLL	Phase-Locked Loop
pMOS	p-channel MOSFET
PPA	Power, Performance, Area
PRNG	Pseudo-Random Number Generator
PSI	Paul Scherrer Institute
PVT	Process, Voltage, Temperature
PWL	PieceWise Linear
RADEF	RADiation Effects Facility
RF	RadioFrequency
RHBD	Radiation Hardening By Design
ROM	Read-Only Memory
RPP	Rectangular ParallelePiped
RTL	Register-Transfer Level
RVT	Regular Voltage Threshold
RWDD	Random-Walk Drift–Diffusion
SAA	South Atlantic Anomaly
SBU	Single-Bit Upset
SCC	SpaceCraft Charging
SCM	Scanning Capacitance Microscopy
SCR	Silicon-Controlled Rectifier
SEB	Single-Event Burnout
SECDED	Single-Error Correction, Double-Error Detection
SEE	Single-Event Effect
SEFI	Single-Event Functional Interrupt
SEGR	Single-Event Gate Rupture
SEL	Single-Event Latch-up
SEMM	Soft Error Monte Carlo Model
SEP	Solar Energetic Particle
SER	Soft-Error Rate
SET	Single-Event Transient
SEU	Single-Event Upset
SIMS	Secondary-Ion Mass Spectroscopy
SoC	System on Chip
SOI	Silicon On Insulator
SPENVIS	SPace ENVironment Information System
SPHD	Single-Port High Density
SPICE	Simulation Program with Integrated Circuit Emphasis

SPL1$	Single-Port Level-One Cache
SQL	Structured Query Language
SRAM	Static Random Access Memory
SRH	Shockley–Read–Hall
SRIM	Stopping and Range of Ions in Matter
STI	Shallow Trench Isolation
TCAD	Technology Computer-Aided Design
TEM	Transmission Electron Microscopy
TIARA	Tool suIte for rAdiation Reliability Assessment
TID	Total Ionizing Dose
TNID	Total Non-Ionizing Dose
TRIUMF	TRI-University Meson Facility
UCL	Université Catholique de Louvain
UML	Unified Modeling Language
UTBB	Ultra-Thin Body and Buried oxide
VDL	Vernier Delay Line
VLSI	Very Large Scale Integration
XML	Extensible Markup Language

Symbols

Radiometry, macroscopic quantities

Symbol	Meaning	Unit
ϕ	Particle flux	$\mathrm{cm}^{-2} \cdot \mathrm{s}^{-1}$
Φ	Particle fluence	cm^{-2}
$N_{\mathrm{part}}, N_{\mathrm{evt}}$	Particle, Single-Event Effect count	1
$A_{\mathrm{emi}}, A_{\mathrm{tgt}}$	Surface areas (emitter, target)	cm^2
LET_{th}	Linear Energy Transfer (e.g. threshold)	$\mathrm{eV} \cdot \mathrm{cm}^{-1}$ ($\mathrm{MeV} \cdot \mathrm{cm}^2/\mathrm{mg}$, $\mathrm{fC}/\mu\mathrm{m}$)
σ_{sat}	Circuit cross section (e.g. saturation)	cm^2
SER	Soft-Error Rate	s^{-1} (day^{-1}, year^{-1})

Energetic particle physics (“charge deposition”)

Symbol	Meaning	Unit
Z, A	Atomic and mass numbers of nucleus	1
E_{ion}	Ion energy	eV
δ	Delta ray energy	eV
E_{eh}	Electron–hole pair creation energy	eV
ρ_{Si}	Medium mass density (e.g. silicon)	$\mathrm{g} \cdot \mathrm{cm}^{-3}$
N_{Si}	Medium atomic density (e.g. silicon)	cm^{-3}
$\sigma_{\mathrm{p}^+-\mathrm{Si}}$	Nuclear cross section (e.g. p^+ in Si)	cm^2 ($\mathrm{barn} = 10^{-24}\ \mathrm{cm}^2$)
λ	Particle mean free path	cm
l, R	Ion path length and range	cm
Q_{dep}	Deposited charge	C

Semiconductor physics ("charge collection")

Symbol	Meaning	Unit
q	Elementary charge	C
$\varepsilon_{\mathrm{Si}}$	Dielectric permittivity (e.g. silicon)	$\mathrm{F} \cdot \mathrm{cm}^{-1}$
V	Electrostatic potential	V
$\mathcal{F}$	Electric field	$\mathrm{V} \cdot \mathrm{cm}^{-1}$
n, p, n_{i}	Carrier density (electron, hole, intrinsic)	cm^{-3}
t_{Si}	SOI film thickness	cm
$N_{\mathrm{d}}, N_{\mathrm{a}}$	Dopant concentration (donor, acceptor)	cm^{-3}
k_{B}	Boltzmann constant	$\mathrm{J} \cdot \mathrm{K}^{-1}$
T	Absolute temperature	K
V_T, V_{bi}	Thermal voltage, built-in voltage	V
$E_{\mathrm{c}}, E_{\mathrm{v}}$	Conduction, valence band energy	eV
$E_{\mathrm{F}}, E_{\mathrm{i}}$	Extrinsic, intrinsic Fermi level	eV
$E_{\mathrm{g}}, qV_{\mathrm{bgn}}$	Energy gap, BandGap Narrowing	eV
μ_n, μ_p, μ_a	Carrier mobility (electron, hole, ambipolar)	$\mathrm{cm}^2 \cdot \mathrm{V}^{-1} \cdot \mathrm{s}^{-1}$
D_n, D_p, D_a	Diffusion coefficient (electron, hole, ambipolar)	$\mathrm{cm}^2 \cdot \mathrm{s}^{-1}$
$\mathcal{J}_n, \mathcal{J}_p$	Current density (electron, hole)	$\mathrm{A} \cdot \mathrm{cm}^{-2}$
R_{Auger}	Carrier recombination rate (e.g. Auger)	$\mathrm{cm}^{-3} \cdot \mathrm{s}^{-1}$
G_{ion}	Carrier generation rate (e.g. ionizing radiation)	$\mathrm{cm}^{-3} \cdot \mathrm{s}^{-1}$
β	Parasitic bipolar gain	1
$\tau_{\mathrm{r}}, \tau_{\mathrm{f}}$	Radiation current rise and fall times	s
Q_{coll}	Collected charge	C

CMOS circuit analysis

Symbol	Meaning	Unit
$V_{\mathrm{d}}, V_{\mathrm{g}}, V_{\mathrm{s}}, V_{\mathrm{b}}$	FET terminal voltages (drain, gate, source, bulk)	V
V_{t}	FET threshold voltage	V
$L, L_{\mathrm{g}}, L_{\mathrm{min}}$	Gate length (e.g. technology's minimum)	cm
W	Transistor width	cm
$V_{\mathrm{dd}}, V_{\mathrm{nom}}$	Supply voltage (e.g. nominal)	V
$I_{\mathrm{on}}, I_{\mathrm{off}}$	Steady-state drain current (saturation, leakage)	A
$I_{\mathrm{d}}, I_{\mathrm{max}}$	Transient drain current (e.g. max. under radiation)	A
R_{on}	Resistance (e.g. FET in 'on' state)	Ω
G_{on}	Conductance (e.g. FET in 'on' state)	sm
C_{gg}	Capacitance (e.g. FET's gate)	F
Q_{crit}	Critical charge	C

Geometry, probability, numerical methods

Symbol	Meaning	Unit
x, y, X, Y	Horizontal coordinates and lengths	cm
z, Z	Vertical coordinate and length	cm
θ, φ	Spherical angles (tilt, roll)	rad
ω	Solid angle	sr
$\boldsymbol{u}, \boldsymbol{n}$	Unit vector, normal vector	1
P	Probability	1
$E(X)$	Expectation value of random variable X	1
f_X	Probability Density Function of X	$[X]^{-1}$
F_X	Cumulative Density Function of X	1
$B(x)$	Bernoulli function	1
J_f	Jacobian matrix of function $f(x)$	$[f] \cdot [x]^{-1}$

Introduction

Arthur:	Camelot!
Sir Galahad:	Camelot!
Lancelot:	Camelot!
Patsy:	It's only a model.

Monty Python and the Holy Grail

A little over a century ago in 1912, Austrian-American physicist Victor Franz Hess flew a hot air balloon to an altitude of 5,300 meters and measured the air ionization rate to be four times higher than at ground level. At the time, it was believed that atmospheric electricity was only caused by radioactive elements from the ground or isotopes of radon in the air. Therefore Hess' observation could only mean that, in his words, *"a radiation of very great penetrating power enters our atmosphere from above"* [17]: he had just found proof for the existence of cosmic rays, and in 1936 he received the Nobel prize in physics for his discovery.

After the first transistors were devised following World War II, and could be manufactured into a monolithic Integrated Circuit (IC) by the end of the 1950s, it was not long until scientists realized that ionizing radiation could perturb their operation. The first malfunctions to ever be observed came in the 1960s, when nuclear testing above ground triggered electronic anomalies—but for obvious reasons, their observers did not rush to claim such discoveries. As for space exploration, it was in the middle of the 1970s that cosmic-ray induced bit flips were first reported, occurring in the electronics of Hughes satellites. Such "soft errors" have been an active area of research ever since, and as a matter of fact, nowadays even at ground level they are a prominent contributor to the failure rates observed in modern chips. This is because decades of continuous miniaturization have raised the transistor packing density to a point where it is possible for a single ion to affect dozens of cells at once, causing the soft-error rates to increase dramatically, if left unmitigated. Now, while on the ground cosmic ray fluxes are very low—owing to the Earth's magnetic shield—in outer space the sheer amount of particles can lead to intolerable error rates, if nothing is done to "harden" circuits against radiation.

For that reason, historically space-bound electronic circuits have been derived starting from commercial technologies, and taking up a few years to secure proper radiation hardening. These development and qualification times have led to the space industry lagging two or three generations behind what was commercially available in terms of performance. Furthermore, rad-hard design strategies have oftentimes been based on uncompromising hardening mechanisms coming at a high cost either in silicon area,

power consumption, or performance. Recent evolutions in the space industry, however, are shifting the practices towards high-end technologies in order to meet commercial expectations, for example with applications targeting broadband communication with low latency. In order to shorten the time-to-market and reduce development costs, this implies relaxing the mission profiles, with shorter satellite lifetimes, oftentimes in Low Earth Orbit (LEO), with environments less harsh than Geosynchronous Earth Orbit (GEO). This leads to the development of less-conservative rad-hard chips, and with bolder risk management setting the cursor in the performance–reliability trade-off. In that context, this thesis contributes to accompanying the space industrial revolution, by improving the understanding and modeling of single-event effects (SEE) of high-end CMOS technologies intended for space applications, with a specific inclination towards high-performance Fully-Depleted Silicon On Insulator (FD-SOI) technology.

To that end, in Chapter 1 a deep-dive review of cosmic ray environments and single-event effects is presented, along with the industrial context pertaining to present-day radiation hardening activities. Particle spectra in the near-Earth environment are characterized, and their effects on CMOS ICs are detailed. Physical processes at play in soft errors are examined, which lead to circuit effects like single-event upsets (SEU) or transients (SET). This thesis being oriented toward SEE modeling, state-of-the-art models and simulation tools are also presented, namely Technology Computer-Aided Design (TCAD) for low-level characterization of the charge transport mechanisms, and Monte Carlo simulators for circuit-level Soft-Error Rate (SER) prediction. CMOS scaling trends are discussed, underscoring how SETs have progressively become a substantial contributor to the overall SER of modern chips. The behavior of modern alternatives to the planar bulk transistor, such as FD-SOI or Fin Field Effect Transistor (FinFET), is then presented. These technologies offer significant improvements in electrical performance, and are shown to exhibit much lower SER than planar bulk counterparts. These mass-market oriented technologies are thus increasingly being adopted for novel commercial space applications such as short LEO missions driven by cost efficiency, where radiation tolerance may be but one among several constraints to contend with. In light of this scientific and industrial context, the research objectives of this work are stated: a first goal is to understand the radiation response of advanced FD-SOI, by means of TCAD simulations on elementary transistors; based on this, one of the main objectives of this work is to establish compact SEE models, so as to enable numerically-efficient and accurate Monte Carlo SER predictions at logic-cell level in FD-SOI. Finally, SER prediction capabilities shall be used in the two main space technologies addressed within this work, namely 65 nm bulk Si, and 28 nm FD-SOI. In complement with irradiation measurement data, Monte Carlo SER simulations must allow to characterize the SEU and SET response of these platforms, and enable practical hardening guidelines and design optimizations for space-bound chips manufactured in modern technologies.

Chapter 2 of this manuscript describes Tool suIte for rAdiation Reliability Assessment (TIARA), the Monte Carlo SER used, maintained, and upgraded, within this PhD. After an overview of the tool's architecture and main features, all of TIARA's physical modules are discussed at length. This includes random sampling of particle distributions according to the radiation environment, energetic particle interaction with silicon, semiconductor processes of carrier transport and collection, and circuit response with calls to an external SPICE solver. TIARA's software layers are also described, with a special emphasis on automated generation of simulation structures in the physical and electrical domains, based on circuit descriptions from the standard CAD flow. The result post-processing module and the simulation harness are also covered, the latter includ-

ing distributed execution via parallel dispatching of particle batches on a clusterized infrastructure. A comparison with other state-of-the-art tools is also provided.

Chapter 3 presents the modeling work carried out in FD-SOI: TCAD simulations on elementary transistors are performed, allowing to delimit the reduced sensitive volumes for radiation-induced charge collection. Parasitic bipolar amplification mechanisms are also studied, and shown to be quite contained in 28 nm FD-SOI. This leads to the development of a first compact SEE model in FD-SOI, accounting for a parasitic bipolar gain as a function of ionizing power, and a position-dependent collection efficiency. This behavioral, bias-independent model is applicable for SEU prediction. A second FD-SOI SEE model is then proposed, based on a one-dimensional drift–diffusion solver running transient transport simulations within the impacted devices. The model is dynamically linked to the electrical solver operated by TIARA, allowing to dynamically capture the circuit feedback on charge collection. The model is confronted to, and proved to be in agreement with mixed-mode 3D TCAD simulations, and therefore it can be applied to both SEU and SET studies.

Finally in Chapter 4, Monte Carlo SER simulation studies are conducted in 65 nm bulk Si and 28 nm FD-SOI; the new FD-SOI collection models are first extensively validated against experimental data on bit-cells, then TIARA is used for predictive simulations on sequential and combinational logic cells. Several circuit design options are investigated in 28 nm FD-SOI, showing how Monte Carlo simulations can allow to complement beam test data, facilitating explorations across the cell libraries, operating conditions, and particle environments. This is especially true in FD-SOI technology, whose native robustness sets very sharp constraints for experimental test plans, that must then browse a restricted set of test conditions. Furthermore, a study of single-event transients in 65 nm bulk Si clock trees is presented. This leads to issuing practical SET guidelines to protect such high fan-out nets that can otherwise be a source of Single-Event Functional Interrupt (SEFI). Finally, on-orbit SER estimations are discussed, demonstrating the benefit of FD-SOI over planar bulk technology in terms of soft-error resilience. Future modeling challenges are also highlighted.

Chapter 1

Background on single-event effects in space electronics

This introductory chapter serves three main goals: to first present some scientific context relevant to the general domain of space radiation effects on electronics, by giving a description of the radiation environments encountered on orbit (Section 1.1) and their impact on an IC (Sections 1.2, 1.3). Then, to establish the state of the art in the more specific field of soft-error simulation that this thesis addresses (Section 1.4). Finally, with all of the above in mind and after reviewing the industrial challenges associated with the space market evolutions, to define the perimeter and objectives of this research work (Section 1.5).

Before we begin, let us delineate the scope of this thesis a bit more: the effects of radiation on circuits can be manifold, as shown in Figure 1.1. The focus of this research is on Single-Event Effects (SEEs), which means studying the effects triggered by a single particle strike on a circuit. Most of the time, the effects are reversible and thus termed "soft errors". Therefore we often use both terms interchangeably, even though single-event "hard errors" do happen and ought not to be neglected—they just were not in the primary scope of this work. Also, while Displacement Damage (DD) i.e. proton-induced defects in the semiconductor lattice, is mostly a concern for photovoltaic applications, Total Ionizing Dose (TID) is by no means an "exotic" effect: radiation-induced charges can *flow* in conductors or semiconductors (leading to SEEs), but in insulating structures they *accumulate*, which can give rise to parametric shifts on the irradiated devices (typically a degradation of the transistors' threshold voltage as oxide charge builds up over the years in space). TID was indeed addressed during this PhD program, but mostly from an engineering perspective, thus in this text we are not reporting contributions on that topic. On a side note, let us mention that DD is sometimes also called Total

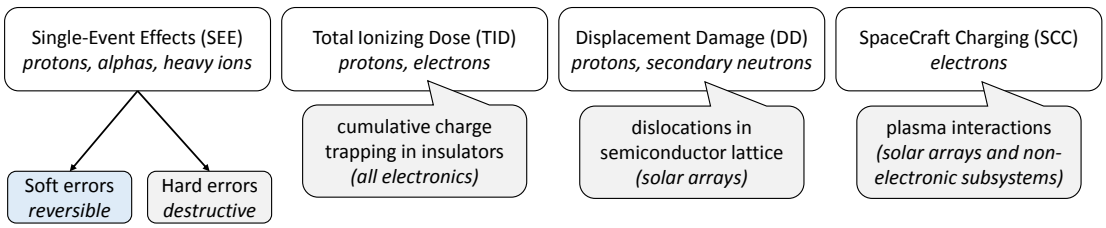

Figure 1.1: Classification of space radiation effects on electronics.

Non-Ionizing Dose (TNID), especially in Europe. Lastly, the interactions with energetic plasmas in orbit may affect many subsystems in a satellite (via arcing discharges or sputtering of surface materials to just name a few effects); in electronics, SpaceCraft Charging (SCC) primarily manifests as an ElectroStatic Discharge (ESD) issue, which is also outside the scope of this thesis.

1.1 Space radiation environments

Cosmic rays are defined as high-energy radiation of extraterrestrial origin. The usage of the word *ray* originates from their discovery in the early twentieth century when they were—wrongly—thought to be of electromagnetic nature; nowadays, the term usually refers to *massive* particles of matter capable of acquiring great energy by accelerating to relativistic velocities. On the other hand, massless photons always travel at the same speed and can only achieve high energy when their frequency is high (i.e. gamma rays or X-rays): the energies involved in gamma bursts (cosmic photons) are several orders of magnitude below those of fermionic cosmic rays.

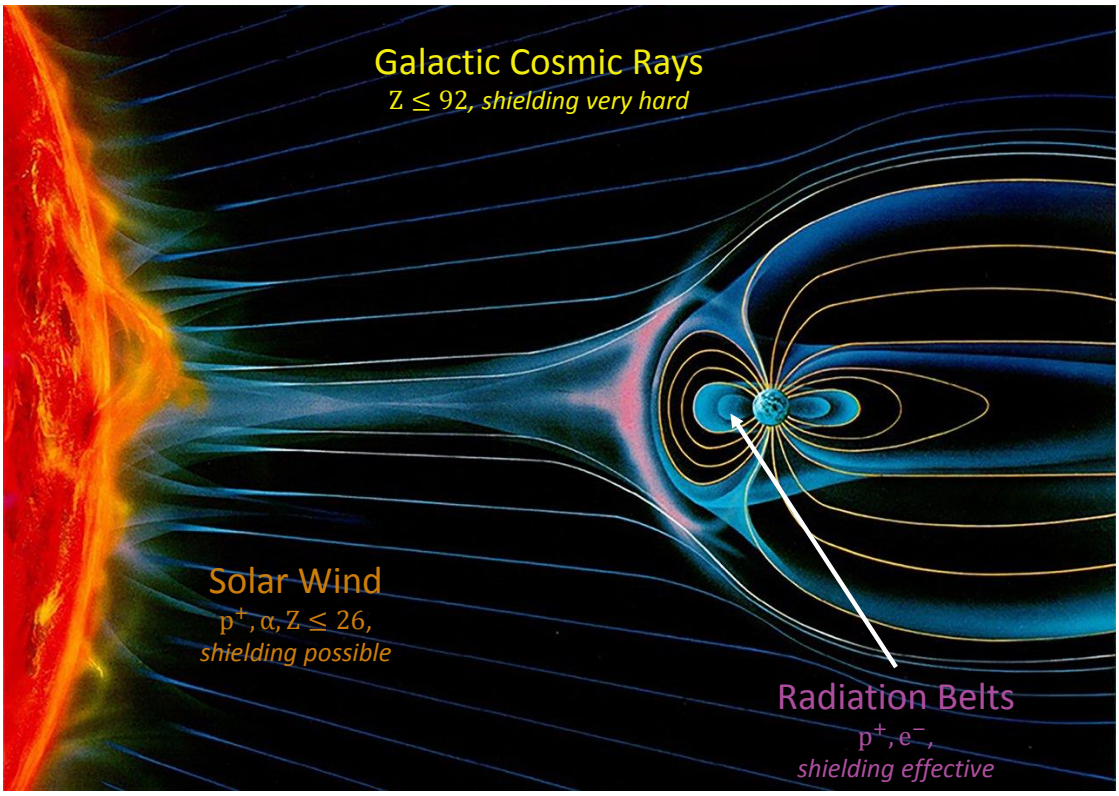

Figure 1.2: An artist rendering of space radiations and their interaction with the Earth's magnetic field. Created by artist K. Endo, courtesy Prof. Yohsuke Kamide, Nikkei Science.

Even though astronauts on the International Space Station (ISS) have reported perceiving occasional "light flashes" when closing their eyes[1], for the most part, cosmic rays are invisible. Yet they are present everywhere and cannot always be stopped with shielding material. Figure 1.2 depicts the main classes of particles encountered in space, more specifically near the Earth as a result of the (partial) protection provided by its

[1]Such cosmic ray visual phenomena originate from the interaction of radiation with the eye, but their exact mechanism is yet uncertain.

magnetosphere. At this point, let us mention that this PhD work is not oriented towards characterizing the space radiation spectrum; we address the *effects* of radiation on semiconductor devices, after the on-orbit environment has been simulated in tools such as Cosmic Ray Effects on Micro-Electronics (CRÈME) from NASA and Vanderbilt University [18], or ESA's SPace ENVironment Information System (SPENVIS) [19]. In the next few paragraphs we will browse through the zoology of cosmic rays and discuss phenomenologically how they are quantified in such simulators. Later on in this thesis this will give us an input point for electronics-related considerations—but we do not challenge environment models *per se*.

1.1.1 Galactic cosmic rays

A Galactic Cosmic Ray (GCR) is an energetic particle from outside the solar system. About 99% of them are nuclei (atoms stripped of their electrons shells), the other 1% comprising mostly of beta particles i.e. solitary electrons (β^-, e^-)—and the remaining fraction is made of stable antimatter such as positrons (β^+, e^+) and antiprotons. Of the bare nuclei, or in other terms ions, by decreasing abundance there are about 90% of hydrogen nuclei i.e. protons ($^1H^+$, p^+), about 9% of helium nuclei also called alpha particles ($^4He^{2+}$, α), and the remaining 1% are High Z and Energy (HZE) ions spread across the entire periodic table of the elements, i.e. with atomic number Z up to 92 for uranium[2]. Note that HZE ions are usually known to the radiation effects community as "heavy" ions, the term HZE being more frequent in astrophysics.

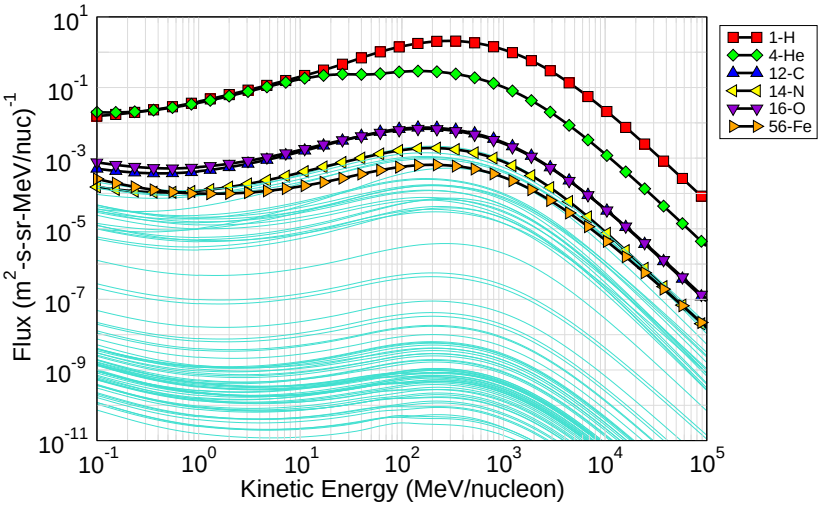

Figure 1.3: Fluxes of ions species encountered on a Geosynchronous Earth Orbit (GEO) (CRÈME simulation). The legendless cyan curves represent nuclei up to uranium.

The particle spectrum for a typical Geosynchronous Earth Orbit (GEO) simulated with CRÈME is shown in Figure 1.3. Note that the quantity on the ordinate axis is an energy- and angle-dependent flux. We shall take this opportunity to introduce some useful quantities, with notations and definitions inspired by recommendations from the

[2] 83 out of the 118 known chemical elements have at least one stable isotope, or stable enough to be older than Earth (4.5 billion years). Three of them are radioactive, the heaviest of which is uranium, hence the "Z=92" limit that we mention for GCRs.

International Commission on Radiation Units and Measurements (ICRU) in [20]: let ϕ (lowercase phi) be the total flux i.e. the number of ions per unit area (e.g. in square centimeters) per unit time (e.g. in seconds). Then $\frac{\partial \phi}{\partial E\, \partial \omega}$ on the figure's y-axis measures the energy and direction dependence of the flux where we denote the energy E e.g. in mega-electronvolts, and the solid angle ω e.g. in steradians. We will thus have:

$$\phi = \iint\limits_{E,\omega} \frac{\partial \phi}{\partial E\, \partial \omega}\, \mathrm{d}E\, \mathrm{d}\omega = \iiint\limits_{E,\theta,\varphi} \frac{\partial \phi}{\partial E\, \partial \theta\, \partial \varphi}\, \mathrm{d}E \cdot \sin\theta\, \mathrm{d}\theta\, \mathrm{d}\varphi \tag{1.1}$$

where θ and φ (a variation on lowercase phi) are the spherical angles, respectively for tilt (or zenith) and roll (or azimuth). The equation above establishes the relationship between a *differential* and an *integral* particle flux. Beside the instantaneous flux, a quantity called *fluence*, which we denote Φ (capital phi), is also used ubiquitously:

$$\Phi = \int_t \phi\, \mathrm{d}t \iff \phi = \frac{\partial \Phi}{\partial t} \tag{1.2}$$

That is, the fluence is the flux integrated over time t, analogous to the *exposure* used mostly in radiometry i.e. for electromagnetic waves, but sometimes also by radiation effects scientists.

Now back to Figure 1.3—as can be seen, protons and alpha particles make up for most of the population of GCRs, but heavier ions are also present, routinely achieving GeV energies. They have relativistic velocities, which means they can hardly be stopped by shielding layers from a satellite, spacecraft, or instrument. Likewise, some of the ions with highest energy-mass are not successfully deflected by the Earth's field, especially around the poles where the magnetic lines reconnect. On a side note, this is why GCRs and their secondary products (after cascading nuclear reactions in the atmosphere) are a prominent cause of SEEs at atmospheric altitudes, namely for avionics or terrestrial applications. Although their intensity is somewhat modulated by the solar activity as we will see right after, this all means that GCRs have a mainly *persistent* nature, as opposed to transient or dynamic—the current observed rates of GCRs (with integral flux in the $1/\mathrm{m}^2/\mathrm{s}$ range) are the result of astrophysical processes occurring on cosmological time scales and that are still widely debated, hence the rather empirical nature of most GCR models implemented in radiation environment codes.

1.1.2 Solar energetic particles

Quite on the contrary, the activity of the Sun can be quite fickle. It is paced by two kinds of events, namely Coronal Mass Ejections (CMEs) and solar flares. The former are large releases of matter and electromagnetic radiation above the Sun's surface, thus CMEs primarily produce a slow-moving electron and proton plasma, albeit with some photons. The latter kind of events, solar flares, consist of a sudden increase in the Sun's brightness, also accompanied by the ejection of clouds of matter: protons, electrons, alpha particles, and HZE ions. Such solar flares and, more rarely, the shock waves associated with CMEs, give rise to the Solar Energetic Particle (SEP) spectrum that we receive from the Sun. Solar particles can be as heavy as iron (Z=26, the heaviest element produced by the stellar nucleosynthesis process), with energy ranges typically lower than a GCR. Consequently, satellite shielding can screen a portion of the solar wind that may harm the on-board electronics. Note that CMEs and solar flares often coincide, but not always.

The overall frequency of these events is modulated by an eleven-year "solar cycle", which has been observed for centuries in the occurrence rate of auroras or in the varying number of sunspots for instance. This periodicity pertains to how the Sun's complex magnetic field evolves over time, and gives rise to so-called solar maxima and minima separated by a half cycle, as can be seen in Figure 1.4: the flux of energetic solar protons can be as low as $100/\mathrm{cm}^2/\mathrm{s}$ and as high as $10^6/\mathrm{cm}^2/\mathrm{s}$. In environment simulation tools typically, SEP models capture these dynamics by providing either short term fluxes (e.g. worst five minutes, worst day, worst week solar flares), or long-term fluences which can be divided by larger time scales to establish averaged-out fluxes. The solar cycle also influences GCRs: the overdensity of plasma induced by CMEs during solar maximum scatters a portion of the cosmic ray background and therefore, the flux of GCRs is *anticorrelated* to solar activity—it is at its highest during solar minimum. Since the temporal variation of the GCR rate is not as harsh as for solar particles, environment models have been historically based on extremal configurations, for example "May 1996" for a GCR worst case which is the second-to-last recorded solar minimum at the time of this thesis.

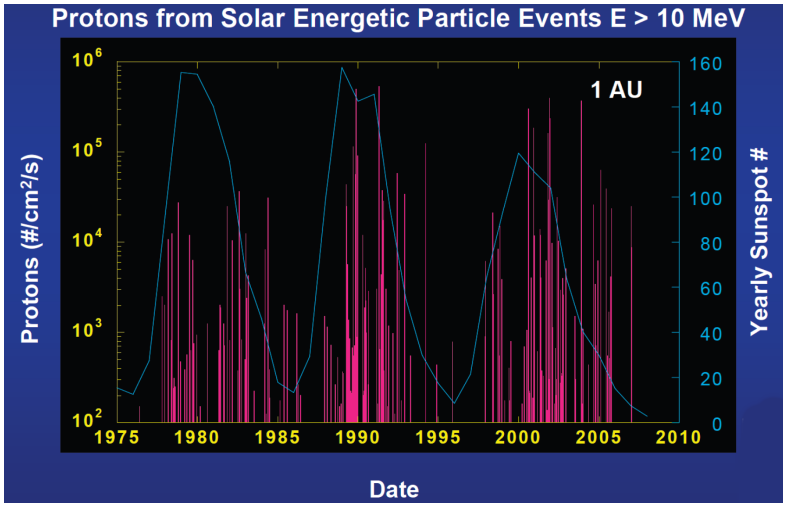

Figure 1.4: Three solar cycles evidenced by the rate of occurrence of large proton events from solar wind, or proton storms, in correlation with the number of sunspots [21].

1.1.3 Trapped particles, radiation belts

Beside deflecting a fair portion of the cosmic rays approaching the Earth, the magnetosphere can also have the adverse effect of *trapping* them inside so-called Van Allen belts, after the name of their American discoverer in the late 1950s. There are two radiation belts around our planet, as represented in Figure 1.5a: an inner belt ranging from 1000 to 6000 km above sea level (i.e. with a radius of 0.2 to 2 Earth radii), and an outer belt of altitude 13 000 to 60 000 km (3–10 Earth radii). They are mainly populated by solar wind and therefore, consist mostly of protons and electrons—the highest concentrations of energetic electrons being found in the outer belt, while the inner belt is mainly composed of protons instead. Thus depending on altitude, fluxes can range from a few dozen to a few thousand protons or electrons per square centimeter per second. As can be seen in Figure 1.5a, the Van Allen belts are of toroidal shape, which directly arises from the Earth's dipole field: owing to the Lorentz force, charged particles travel a helicoid following the field lines that reconnect at the poles, and drift around

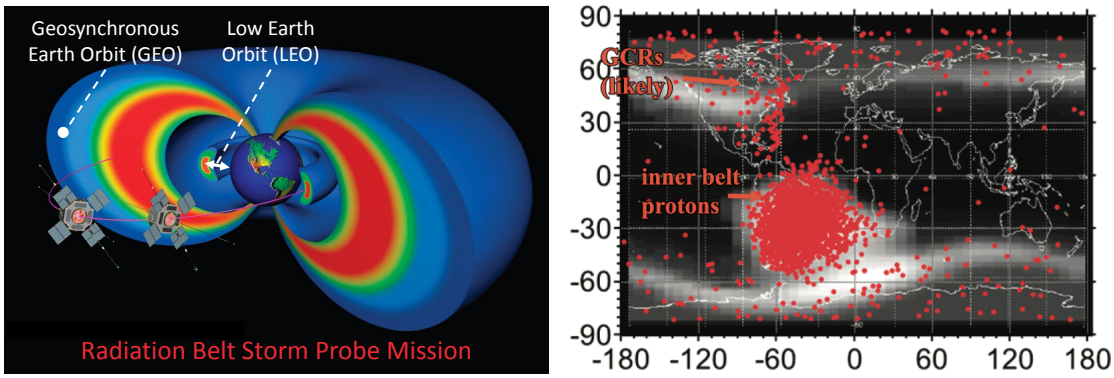

(a) (b)

Figure 1.5: (a) The Van Allen radiation belts are areas with high concentrations of protons and electrons trapped by the magnetosphere (image from John Hopkins University, Applied Physics Labs) – (b) The South Atlantic Anomaly (SAA) is a region where the inner belt dips closer to the Earth, leading to higher failure rates in Low Earth Orbit (LEO), as shown here by the SEEs recorded on the SAMPEX mission (the shading is logarithmic intensity of measured protons and electrons above 0.5 MeV) [22].

the Earth in opposite directions depending on their charge. Note that the outer belt is of more variable shape than the inner belt, being more affected by solar wind which makes it contract or expand. A peculiarity of the inner belt is that, because the Earth's magnetic axis is tilted 11° and offset by 500 km from its rotational axis, in a region called the South Atlantic Anomaly (SAA) the inner belt declines to lower altitudes and causes "unusually-high" particle fluxes. This directly translates to increased error rates for satellites in Low Earth Orbit (LEO), as illustrated in Figure 1.5b with the bus retry errors recorded on NASA's SAMPEX mission [22].

Although the kinetic energies carried by trapped particles can be quite high (typically hundreds of keV for electrons and dozens of MeV for protons), they are generally lower than those of SEPs and GCRs and therefore, shielding material can be quite efficient at stopping such particles. Trapped electrons in particular, are only of concern if very thin (e.g. submillimiter) shielding is employed, or for unshielded parts such as solar panels or antennae standing directly in outer space. Besides, as will be discussed later on, their very low ionizing power makes it extremely unlikely for electrons to trigger soft errors, which rather makes them an issue mainly in terms of total dose, either ionizing (TID) or non-ionizing (DD or TNID). Put another way, the radiation hazard from particles harbored by the Van Allen belts lies in their abundance more than their intrinsic individual effects.

1.2 Physics of single-event effects

Now that we have reviewed the variety of cosmic rays that space systems can experience, we move on to present the threats posed by such radiation. This section will introduce the fundamental notions at play in SEEs, by giving an overview of the physical mechanisms involved in the response of an IC to incoming ionizing particles.

Single-event phenomena in digital ICs can be broken down in four main processes, as summarized in Table 1.1: first the initial generation of carriers by the incoming radiation, then the transport and collection of these charges inside the struck device(s) followed

by the circuit's response to these currents, and finally the response at system level, if any. As shown in the table, these conceptual layers are in fact not hermetic, since the spatial and temporal scales involved can overlap significantly. Mathematically speaking, this means that the underlying equations need to be coupled if one wishes to develop a self-consistent simulation approach. The work in this thesis is clearly focused on the second and third bricks, namely the study of how radiation-induced carriers travel inside silicon and get collected at circuit nodes, thereby affecting the circuit's state. Electrical state changes may then alter charge transport in return, and so on and so forth.

	Time scale	**Spatial scale**	**Physics**	**Tools**
Charge deposition	0.1 fs – 1 ps	10 nm – 1 µm	Particle scattering, nuclear physics	Monte Carlo nuclear codes
Device response	1 ps – 10 ns	1 µm	Semiconductor physics	Device simulators ("TCAD")
Circuit response	1 ps – 1 µs	1 – 100 µm	Circuit theory	Analog circuit simulators ("SPICE")
System response	1 ns – ∞			Digital circuit, logic simulators

Table 1.1: Description layers for single-event effects. Adapted from [23].

1.2.1 Particle interaction with silicon – initial charge generation

The topic of particle–matter interaction could take up an entire textbook in its own right, so here we shall simply give a notion of the main concepts needed for soft-error scientists and engineers. Thus we only describe the effect on silicon of atomic nuclei (ions) and their hadronic constituents (neutrons and protons), and choose to leave aside other subatomic particles.

Interaction of neutrons with matter

While neutrons are not present in space (since they originate from cosmic ray interactions with the Earth's atmosphere), their interaction with matter can first be described, later to be transposed to understand proton interaction, as will be discussed afterwards. When a neutron penetrates a certain target, since it has no electric charge it does not directly ionize the targeted medium. Instead, it can either be *scattered* by the target nuclei or be *captured* by them, as depicted in Figure 1.6a:

- Scattering of a neutron by an atom from the silicon lattice can either be elastic or inelastic. In the former case, the total kinetic energy is conserved, and the neutron is deflected from its path as it transfers some of its energy to the silicon atom. In the latter case of inelastic scattering, the target nucleus rearranges its internal state to one of higher energy, and the total kinetic energy is not conserved. But in both cases, there is only one outgoing neutron and the nature of the recoil nucleus is left unmodified.

- When the incident neutron gets captured by an atom of the target, the spectrum of possible outcome is wider; after it has absorbed the impinging neutron, the nucleus can get rid of excess protons or neutrons, it can undergo de-excitation by emitting a γ-ray, or it may even split in medium-sized fragments when the energies are high enough to trigger nuclear fission. Put briefly, this is when actual nuclear reactions occur, and predicting the resulting secondaries is the topic of nuclear physics. Visually speaking, a histogram plot of the products' atomic number Z gives a "two-hump" curve for a somewhat bimodal distribution, due to the numerous lightweight spallation fragments accompanied by heavy recoil nuclei, and the less abundant mid-sized fission products.

Such nuclear interactions are addressed by Monte Carlo transport codes like GEometry ANd Tracking (Geant4), Monte Carlo N-Particle (MCNP), FLUktuierende KAskade (FLUKA) [1, 2, 24] which statistically simulate possible trajectories based on the interaction probabilities. These probabilities are measured by the interaction *cross section* of the projectile (e.g. a neutron) against the target (e.g. silicon), which is the *effective* surface of a target atom seen from the projectile. As shown in Figure 1.6b, $\sigma_{\mathrm{n-Si}}$ the total interaction cross section of neutrons against silicon is typically of one barn or 10^{-24} cm^2. Given the atomic density of silicon $N_{\mathrm{Si}} = 5 \times 10^{22}$ cm^{-3}, this means that the neutron's mean free path in silicon is

$$\lambda = (\sigma_{\mathrm{n-Si}} \cdot N_{\mathrm{Si}})^{-1} \approx 20\mathrm{cm}. \tag{1.3}$$

Nuclear interactions of neutrons with silicon are thus quite unlikely—however they should not be overlooked, because the subsequent nuclei are precisely those responsible for ionization: neutrons are said to be *indirectly ionizing*.

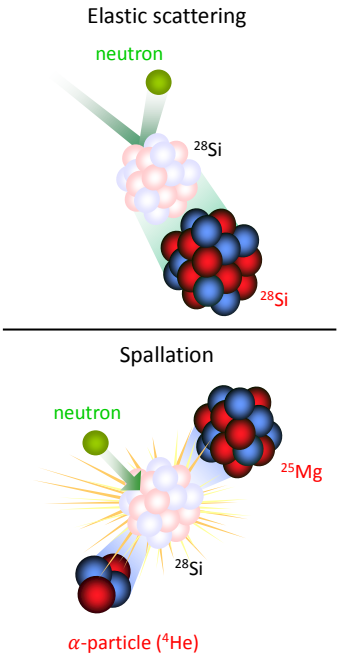

(a)

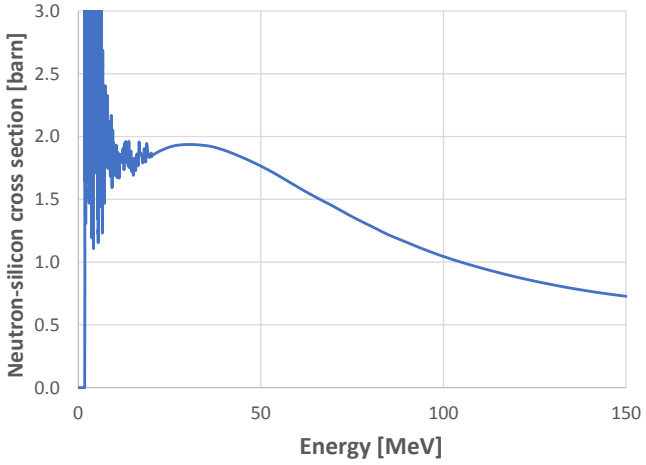

(b)

Figure 1.6: Interaction of neutrons with matter: (a) Examples of neutron–silicon events; elastic scattering, and ${}^{28}_{14}\mathrm{Si} + \mathrm{n} \rightarrow {}^{25}_{12}\mathrm{Mg} + \alpha$, the first reaction channel with 2.75 MeV threshold energy [25] – (b) Total n–Si interaction cross section, from the ENDF/B-VII nuclear database [26].

Interaction of ions with matter

Now, when an atomic nucleus (i.e. an ion of positive charge) travels through matter, it is mostly subject to Coulomb's interaction with electrons from the target material—because, as seen previously, interaction with nuclei from the target is quite unlikely. Therefore, a heavy ion essentially propagates in a straight line since the energy it loses by knocking off the target's electrons is modest. This energy transfer does however slow it down, and eventually the ion will stop after having produced a long (micrometric) and thin (a few dozen nanometer-wide) track of so-called *delta rays*, i.e. secondary electrons having gained enough kinetic energy to themselves ionize the medium. Once again, the processes involved are quite intricate:

- As the ion strips electrons off the target's atoms, it creates a local polarization field which further participates in the slowing-down process. This feedback is especially important at medium energies (typically a few MeV) around the maximum energy transfer, and is studied within the theory of the complex-valued dielectric function governed by Poisson's equation.
- At high energies on the other hand, the electronic slow-down is derived through Bethe–Bloch's quantum theory, which in its non-relativistic version states that the energy transferred per unit length is proportional to the ion's atomic number squared—a remnant from Fermi's golden rule—divided by its energy: unsurprisingly, the heavier the ion, the greater its interaction with the medium, but *the higher the energy, the lower the energy transfer.*
- Finally, at low energies, what we said before about nuclear interactions is not entirely verified: the ion does in fact undergo scattering by the target nuclei, hence a significant fraction of the energy transfer is nuclear instead of electronic and the ion does not travel a purely straight line.

All the complexity of the above is summarized in the concept of Linear Energy Transfer (LET) defined as:

$$LET = (1/\rho_{\text{tgt}}) \cdot \frac{-\partial E_{\text{ion}}}{\partial l} \tag{1.4}$$

where l is the distance that the ion travels along its propagation axis. The parentheses around $1/\rho_{\text{tgt}}$ mean that sometimes the LET is directly expressed as an energy loss per unit length (e.g. in MeV/cm) and sometimes it is normalized by the target's mass density, typically yielding units of $\text{MeV}\cdot\text{cm}^2/\text{mg}$. This normalization is justified by the fact that the mass stopping power, to a first approximation, does not depend on density: an ion traveling through a gas of 1 atm pressure loses about twice as much energy per unit length as it does in the same gas under 0.5 atm—and conversely its range will be about half because the gas is perceived twice "thicker".

It took the scientific community a fair portion of the twentieth century's second half to study, but nowadays the LET of all ions in the periodic table can be predicted within ± 20% accuracy [27], against essentially "any" target. Using the vastly-known and widely-adopted program Stopping and Range of Ions in Matter (SRIM) [6], curves of stopping power (another term for LET) can be obtained as shown in Figure 1.7 for several ion species. As mentioned earlier, at high energies the LET is a decreasing function of energy. At medium energies it goes to a maximum called the Bragg peak and at energies below it becomes dominated by nuclear stopping (see the silicon ion curves in green).

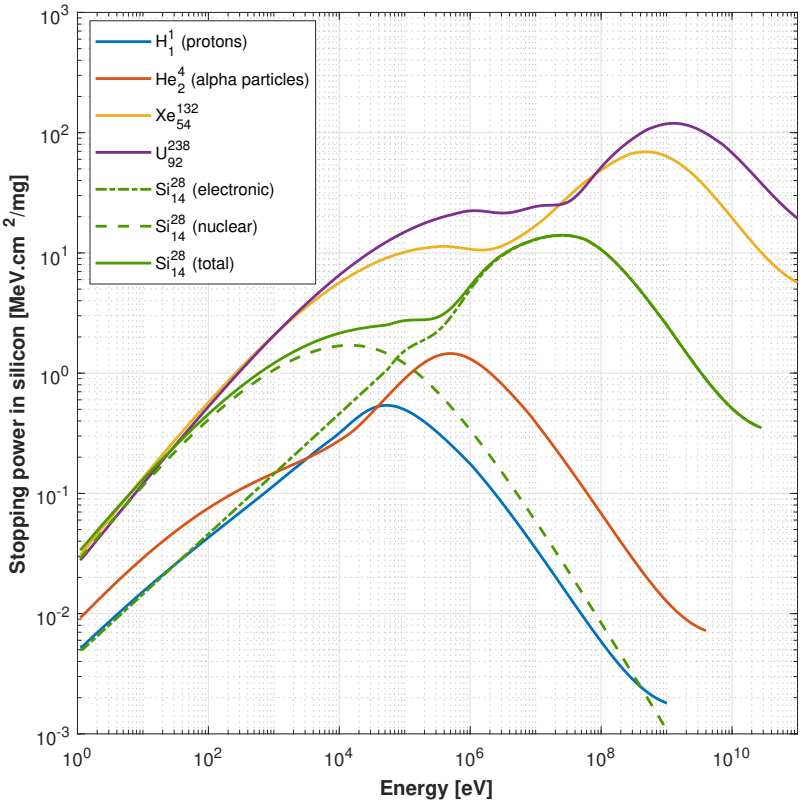

Figure 1.7: Stopping power of several ion species in silicon.

When the target material is a semiconductor, ionization by the δ-rays is described as the promotion of electrons from the valence to the conduction band. In other terms, as shown in Figure 1.8, secondary electrons leave a trail of electron–hole pairs in their wake until they thermalize—from several keV down to $k_\mathrm{B}T = 26$ meV at room temperature, having traveled from a few nanometers to a few microns, in some hundreds of femtoseconds [28]. The average energy needed to create an electron–hole pair is a complex quantity to estimate: intuitively, its order of magnitude is given by the semiconductor's energy gap, but more generally it is determined by its full band structure—several empirical or analytical formulas do exist, see for instance [29]. In silicon, $E_\mathrm{eh} = 3.6$ eV, which can be compared with its bandgap $E_\mathrm{g} = 1.1$ eV: it takes about three times more energy for a δ-ray to create a pair, than for a photon (assisted by a phonon to overcome silicon's indirect bandgap). Knowing that the electron charge is $q = 1.602 \times 10^{-19}$ C and that the density of silicon is $\rho = 2.32$ g/cm^3, multiplying the "mass-LET" by $\rho \cdot q/E_\mathrm{eh}$ we deduce the crucial relationship:

$$1 \text{ MeV} \cdot \text{cm}^2/\text{mg} \iff 0.232 \text{ MeV}/\mu\text{m} \iff 10.3 \text{ fC}/\mu\text{m in silicon} \tag{1.5}$$

Beyond unit conversions, these three variations on the concept of LET all serve a specific purpose: as previously mentioned, the mass-LET allows to abstract the nature of the target to some extent, and furthermore it can directly yield a total ionizing or non-ionizing dose (TID or TNID in energy per unit mass), when multiplied by the fluence of impinging particles; also on a practical note in the SEE field, the mass-LET ranges on an intuitive 1–100 scale, the Bragg peak or maximum LET being in those units close to the atomic number Z. Now the "energy-LET" is useful for energy deposition studies, or high-energy transport calculations; finally, the "charge-LET" is the meaningful quantity

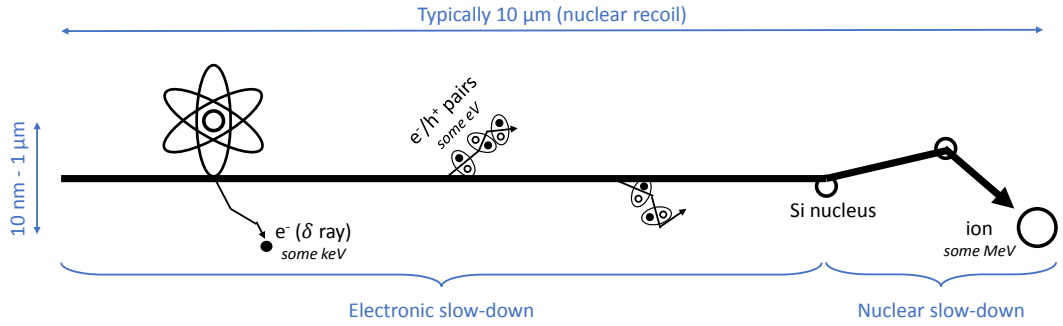

Figure 1.8: Charge deposition by an ion traversing silicon: delta rays are produced which induce electron–hole pairs in the vicinity of the ion track.

for semiconductor physics, or carrier transport calculations, and thus to analyze circuit effects. Note that the three terms we use here are for the sake of clear presentation, and not standard definitions for various LET conventions.

By integrating the reciprocal of the LET over the energy, one may also obtain $R(E_0)$ the range that an ion starting with energy E_0 travels until it stops:

$$R(E_0) = \int_{E_0}^{0} \frac{\partial l}{\partial E} \cdot \mathrm{d}E = \int_{0}^{E_0} \frac{1}{LET} \cdot \mathrm{d}E \tag{1.6}$$

As is made evident by Figure 1.9, ranges of energetic ions can exceed typical satellite shielding thicknesses of a few millimeters—which at this point gives a perhaps more quantitative answer to the layman's question *"why can't you just wrap your chips in aluminum foil altogether?"*; moderate shielding can even cause more harm than good depending on the ion energy, since some secondary nuclear products created occasionally can end up being of higher LET than the primary, which directly increases the likelihood of SEEs, as will be highlighted later on. Note that in the figure the target material is silicon, but stopping powers and ranges in a typical shielding material such as aluminum (next to silicon in the periodic table) would not differ too greatly.

Summary and discussion on protons and electrons

At this point, let us summarize what has been discussed so far for this "initial charge generation" process: when a neutron penetrates a given target, it interacts via nuclear reactions, thereby producing secondary products that carry a charge (light spallation fragments, heavy recoil nuclei, or mid-sized fission fragments). These mechanisms are studied with nuclear physics and simulated in nuclear transport codes. Then when an ion travels through matter, the primary mode of interaction is direct ionization, i.e. the charged nucleus strips off secondary electrons from the target. Delta rays can themselves ionize the medium, inducing a very dense column of plasma whose typical radius is a few dozen nanometers, established in about one picosecond. These processes can also be simulated in Monte Carlo nuclear tools especially if one wishes to model the ion track structure precisely (see for instance [30] and our discussion in Appendix B), but usually the key metric used is the LET given by SRIM.

The interaction of protons with matter lies somewhere in between: at high energies (typically above 30 MeV) their LET is negligible and they can only ionize silicon indirectly, i.e. by triggering nuclear reactions just like high-energy neutrons. As a matter

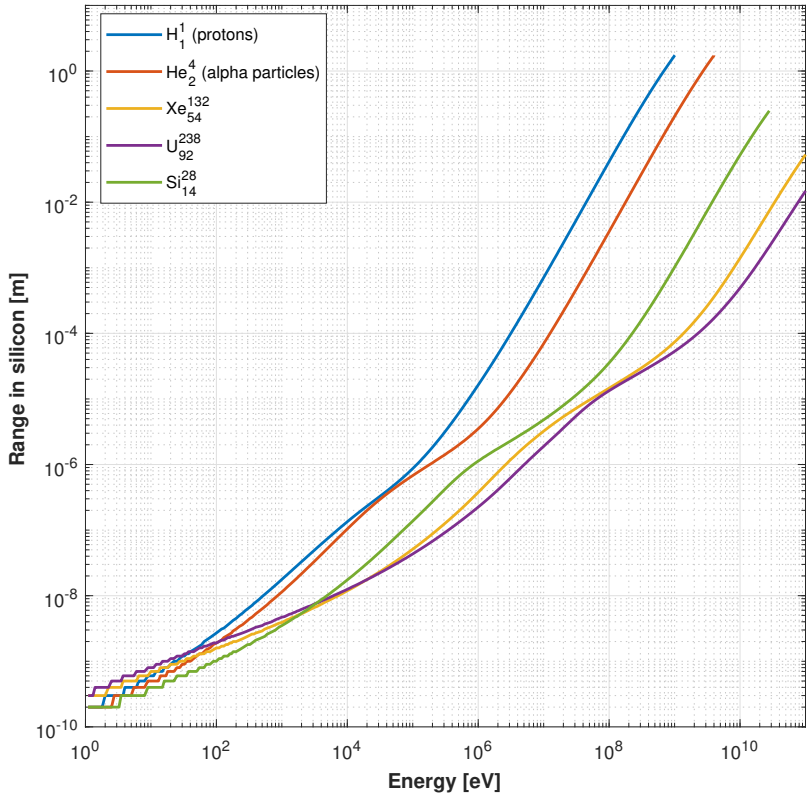

Figure 1.9: Ranges in silicon of several ion species.

of fact, above 50 MeV the nuclear cross sections of neutrons and protons versus silicon become identical, and their effects on chips are similar. At low energies however (below 10 MeV approximately), protons mainly interact via direct ionization. The balance between all these mechanisms will be discussed further at length in Section 4.3 on page 121 when we integrate the response of a circuit over an entire spectrum of energies.

Finally, note that in our description of particle interaction with matter we deliberately did not treat subatomic particles other than neutrons and protons. For space-bound electronics, the main leftovers are electrons, whose erratic motion differs too greatly from that of ions. At low energies collisional losses with other electrons and nuclei dominate, much like for ions but the difference being that the trajectory of an electron is greatly modified when it collides another electron; and at high energies, the stopping power of the electron is dictated by its Brehmsstrahlung. However electrons are only a second-order concern for soft errors since their "equivalent LET" is decades below that of ions, and even then they are stopped by just a few millimeters of shielding.

1.2.2 Semiconductor response – charge transport and collection

Moving one step further, we shall now describe the processes at play in the response of silicon to the charge carriers deposited by the incoming radiation. As a reminder, the previous step of "initial charge creation" involved an ion with energy in the MeV range inducing secondary electrons whose energies lie in the keV range initially, which themselves create electron–hole pairs of no more than a few eVs. The fate of this plasma column thus happens at lower energies usually not handled by nuclear codes, and is now governed by semiconductor physics with larger time and spatial scales (see again

Table 1.1 on page 11). These dynamics are captured by Technology Computer-Aided Design (TCAD) tools such as Sentaurus from Synopsys [31] or Silvaco [32], to just name a few. These software suites allow to model the crucial steps of semiconductor manufacturing process, and then simulate the electrical behavior of the virtual devices obtained. Only a few devices can be simulated in the "physical" domain (i.e. on a three-dimensional grid, under transient excitation for soft errors), due to the computational burden of solving the electrostatic (Poisson's) and carrier transport equations self-consistently. But without resorting to numerous equations, in the next paragraphs we will try to shed light on the transport and collection mechanisms that radiation-induced carriers undergo; in other terms we want to show how charges deposited in silicon can manifest as parasitic currents.

Ambipolar transport under weak field

Current flow in semiconductors is typically understood as comprised of two parts: a *drift* motion of charge carriers in an electric field, plus a *diffusive* part that tends to even out the carrier concentrations—further mathematical details on transport models will be given in Section 1.4.1. Consider a neutral region of silicon to which an external bias is applied, giving rise to a certain electric field. If a charged particle strikes this portion of silicon, the electron–hole pairs induced by the radiation will try to drift in opposite directions due to their opposing charges. But doing so, they disturb the equilibrium neutrality and give rise to a strong internal electric field that will in fact restore neutrality, drawing the electrons and holes back together. This is a phenomenon called *ambipolar transport*, whereby electrons and holes drift and diffuse *unseparated*, with identical mobilities and diffusivities: assuming the excess carrier densities are equal at all time and place, the ambipolar coefficients can be derived to be harmonic means of the coefficients for electrons and holes alone, weighted by their densities.

When transport is truly ambipolar, no net current arises since the $+q$ and $-q$ charges of the carriers balance each other out. In that case, no effects will be seen at circuit level and the electronics will be unaffected by the radiation. On the other hand, if the applied bias is large enough, the external field may overcome the internal restoring fields inside the electron–hole plasma: this is when charge separation occurs. Orders of magnitude for the minimum field required can be estimated by simple means[3], and indicate that it usually takes the electric field of a *reverse-biased p-n junction* to separate the electron–hole pairs of a typical ion track. Put the other way, in forward-biased junctions or silicon regions of constant doping, oftentimes the electric field is not enough to trigger charge separation. This leaves the electron–hole pairs free to diffuse on nanosecond time scales and micrometric spatial scales (related by $L^2 \approx D_a \tau$ where the ambipolar diffusivity D_a is a few cm^2/s), until their densities vanish either by pure diffusion or by recombination.

[3]The characteristic *distance* over which charge separation can occur is given by the Debye length $L_\mathrm{D} = \sqrt{\varepsilon k_\mathrm{B} T / nq^2}$ where ε is the dielectric permittivity and n is the carrier density. It is found by scaling Poisson's equation $\nabla \cdot \mathcal{F} = \rho/\varepsilon$ given the Maxwell–Boltzmann expressions for carrier densities at thermal equilibrium ($\mathcal{F}$ is the electric field and ρ is the net charge density). The characteristic *time* over which charge neutrality is restored is the dielectric relaxation time $\tau_\mathrm{d} = \varepsilon/\gamma = \varepsilon/nq\mu$ (conductivity γ, carrier mobility μ). It is found by substituting $\mathcal{J} = \gamma \cdot \mathcal{F}$ (current density $\mathcal{J}$ is purely conductive assuming drift motion starts to dominate) into Poisson's equation and a simplified continuity equation $\nabla \cdot \mathcal{J} = \partial\rho/\partial t$. At the separation threshold, the carrier velocity is $\mu \cdot \mathcal{F}_\mathrm{th} = L_\mathrm{D}/\tau_\mathrm{d} \implies \mathcal{F}_\mathrm{th} = \sqrt{nk_\mathrm{B}T/\varepsilon}$. Putting the numbers in for the core of an ion track of density $n \approx 10^{18-20}$ cm^{-3} gives a minimum separating field of 10^{5-6} V/cm. Such a field can be reached over the 1 V drop of a p-n junction whose depletion width is 100 nm, but cannot arise in the silicon substrate from an IR drop of 100 mV across a 1 mm chip.

Charge separation and funneling in strong fields

On the other hand, when carriers are injected within an intense electric field, transport can no longer be ambipolar and charges are separated. It was observed early in the radiation-effects community that an order of magnitude for the net charge collected by a reverse-biased p-n junction can be obtained with:

$$Q_{\text{coll}} \approx LET \cdot l_{\text{depl}} \tag{1.7}$$

where l_{depl} is the length of the ion track that is intercepted by the junction's space charge region. At this point, let us mention that, since the depletion width w_{depl} of the junction is a function of the applied bias V_{a}, quite unsurprisingly the collected charge will depend on voltage—at least for an "unrestrained" *bulk* junction i.e. one whose space charge region is not confined by insulators (in Silicon On Insulator (SOI) technology things will be different). For an abrupt junction for instance, $w_{\text{depl}} \propto \sqrt{V_{\text{bi}} + V_{\text{a}}}$ where V_{bi} is the junction's built-in voltage. Different power laws can be obtained for other analytical doping profiles, but the central point is that typically, the collected charge increases less-than-linearly with voltage. For orders of magnitude, remembering (1.5) on page 14, a 100 nm-wide depletion region struck at normal incidence by an α-ray with an LET of 1 MeV$\cdot$cm^2/mg collects an approximate charge of 10.3 fC/μm $\times$ 0.1 μm $\approx$ 1 fC.

Subsequent works [33, 34] noticed however that Q_{coll} can in fact be greater than $LET \cdot l$, even though the estimate seemed to imply a "maximum" (100%) collection efficiency. Careful analysis showed that the plasma column deposited by the radiation can be so dense as to neutralize the junction's depletion region, thereby screening the junction field and pushing the equipotential surfaces deep into the silicon substrate, as illustrated in Figure 1.10. Under this modified electric field, majority carriers are pushed away from the ion track (on the illustration, holes will move radially), while minority carriers remain confined within the track core. They then undergo drift motion (on the illustration, electrons will move longitudinally, upwards), eventually to be collected at the junction electrode.

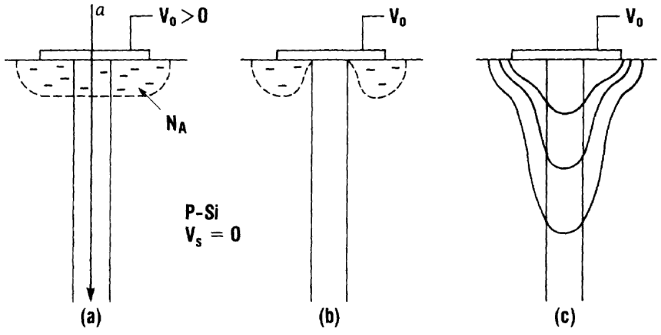

Figure 1.10: Schematic of charge funneling mechanism indicating in (a) an α-particle strike through an n$^+$-p junction and associated well, in (b) depletion layer being neutralized by the plasma column, and in (c) equipotential lines extended down from original junction along particle track. Taken from [33].

The net result is that carriers are collected on a depth greater than the depletion width. This phenomenon was coined "charge funneling" and occurs with quick characteristic times, typically a few picoseconds. The junction field then restores to its original map, starting from the edges of the plasma column where the carrier concentration is close to the background doping density, and working its way towards the center of the

track—which is helped by the fact that the track core dilutes via ambipolar diffusion: the plasma seeks to remain charge neutral at all times, and as long as it is too dense the field cannot separate the carriers (see note 3 on page 17). Analytical models were developed to describe the funneling phenomenon. Without citing the formulas proposed, let us mention that the effective charge collection depth is an increasing function of the LET, which stems from the fact that very dense plasmas are more efficient at screening external fields, causing the equipotential surfaces to extend deeper into the substrate.

On a transition toward circuit-level analyses right after, we may also derive an estimate for the magnitude of the radiation current; considering a *Gedankenexperiment* in which the initial charge cylinder of linear density LET is sucked at a certain velocity v under electric field $\mathcal{F}$, a rule-of-thumb calculation can provide the maximum radiation current as:

$$I_{\max} \approx LET \cdot v = LET \cdot \mu\mathcal{F} \tag{1.8}$$

where μ is the carrier mobility. In practice, this means that typical radiation currents will lie in the milliampere range for particles directly striking a depletion region[4]. Now, whether the circuit will be sensitive to $I_{\max}$ rather than Q_{coll} is a matter of (complex) impedance: if the affected node is of mostly resistive nature (resistance R or conductance G), then the voltage drop (hence the signal perturbation) will be given by $RI_{\max}$; on the other hand, when attacking a node dominated by capacitance C, the voltage drop will scale as Q_{coll}/C. The general case will depend on a total impedance such as $1/(G+jC\omega)$, which we discuss in the next section for CMOS logic. Note how in (1.7) and (1.8), the LET is the governing parameter for the event magnitude: a common assumption in SEE analyses is to consider that an ion is entirely defined by its LET, which is valid provided the sensitive depletion region intercepts the entire ion track diameter, and if across this sensitive volume the ionizing power is roughly constant.

Radiation response of an elementary MOSFET

With the above physics in mind, we can draw the bigger picture for a radiation strike on a semiconductor device such as the MOSFET portrayed in Figure 1.11. When used for digital CMOS logic, statically the nMOS is usually in one of the following two states:

- When the gate is turned on ($V_{\text{g}} = V_{\text{dd}}$ the supply voltage), the transistor is often used to "pass a zero" ($V_{\text{d}} = V_{\text{s}} = V_{\text{b}} = 0$). Thus there are no significant potential gradients for charge separation (except in the thin inversion layer of the FET).
- When the gate is turned off ($V_{\text{g}} = 0$), the nMOS is blocked: it acts as a large resistor and therefore the drain and source voltages are driven by the outer circuitry. In most CMOS logic cases, we will have $V_{\text{ds}} = V_{\text{d}} - V_{\text{s}} \propto V_{\text{dd}}$ depending on the number of off-state nFETs stacked in series—because V_{dd} is reached at the output of some on-state pMOS further away in the logical network. Figure 1.11 represents the one-nMOS case, with dotted lines delineating the depletion regions in active silicon. As illustrated, the largest depletion region is found below the drain electrode with a strong *vertical* electric field arising from the finite drain–bulk voltage, hence the aforementioned collection mechanisms pertaining to reverse-biased p-n junctions will occur: Figure 1.12 depicts a typical current waveform resulting from a radiation

[4]Note that (1.8) does not imply that $I_{\max}$ is proportional to either the voltage or the LET: electric field $\mathcal{F}$ in a junction is not linear w.r.t. voltage, and mobility μ is itself a strong function of LET due to carrier-carrier scattering effects.

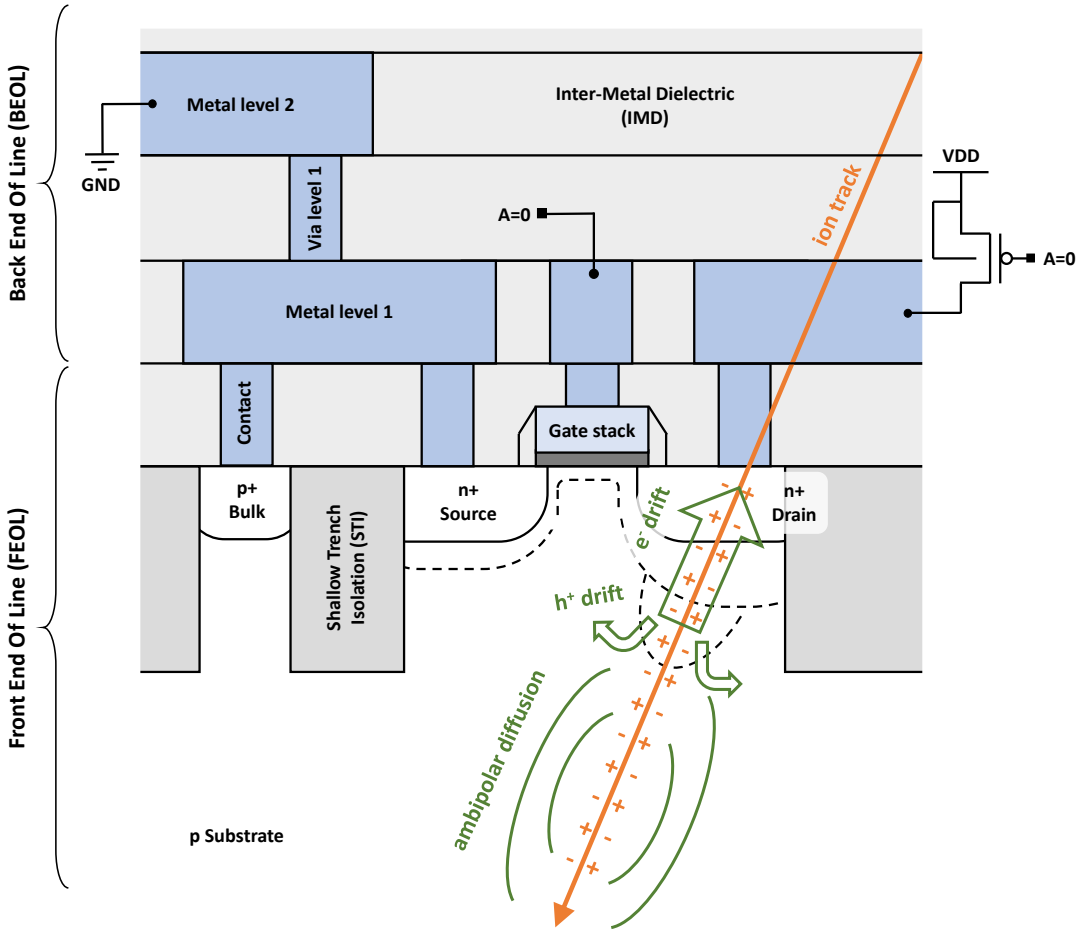

Figure 1.11: Illustration of the charge transport mechanisms in an nMOS subject to a radiation impact, with fast collection from the drain junction's electric field (in dashed lines) and slow ambipolar diffusion. The main technology process features are also represented, roughly to scale for a deep submicron technology. Silicon is in white, conductors in shades of blue depending on conductivity, and insulators in shades of gray depending on dielectric constant.

strike on an nMOS drain (with positive orientation for current flowing into the drain): minority carriers i.e. electrons are directly collected by the junction field (quick drift current with a more or less pronounced funneling), and later on the current exhibits a diffusive tail due to carriers slowly reaching the depletion region via ambipolar diffusion followed by charge separation.

All of this exemplifies the fundamental concepts of radiation-induced charge transport and collection in CMOS technology. Of course, what we said about an nMOS can be switched to discuss pMOS transistors. The "bulk" terminal formed by the p-type substrate tied to GND becomes the n-well tied to VDD in which the pMOS lies; the drain and source electrodes are now p^+ implants. In the 'on' state we have $V_g = 0$ and $V_d = V_s = V_b = V_{dd}$ with no significant electric field arising, and conversely in the 'off' state we have a strong vertical electric field at the drain junction because of the finite $|V_{db}|$. Holes will be collected instead of electrons but in both cases, radiation-induced currents are of positive sign from "plus" to "minus", simply following the field direction—which in most bias cases corresponds to holes (resp. electrons) being pulled towards p-type (resp. n-type) regions.

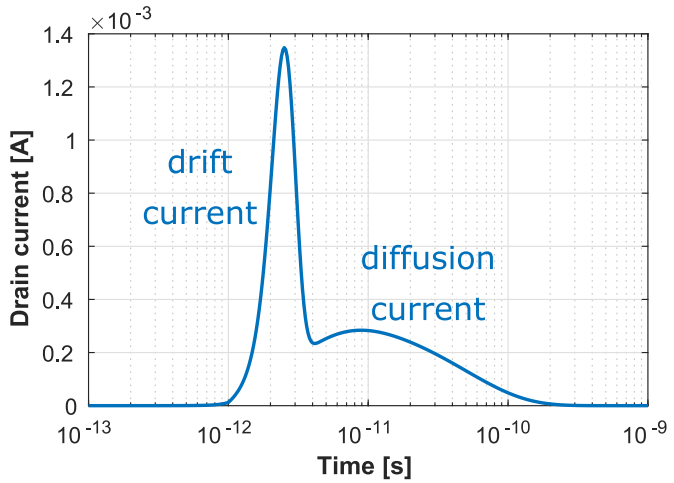

Figure 1.12: Typical radiation-induced current in an off-state transistor, with prompt collection from the drain junction and longer-lasting diffusion tail.

Finally, note that in this discussion we ignored the charges deposited outside the active silicon. This assumption is justified by two facts:

- The conductivity of Back End Of Line (BEOL) metals is orders of magnitude above that of semiconductors, which means charge neutrality is achieved extremely quickly (see note 3 on page 17) and we need not worry about those areas. Note that a further approximation that is sometimes used is to also neglect charges deposited in the heavily-doped implants ("decent" conductors), because the numerous recombination centers provided by the impurities lead to short lifetimes for the excess carriers (typically a few dozen picoseconds).
- Insulators, by definition, are materials in which charges can only move very slowly. This is true of Shallow Trench Isolation (STI) regions, transistor gates (which additionally are very thin, thus not intercepting many charges), and BEOL dielectrics. On top of that, insulators have a larger energy gap than silicon, thus leading to a higher electron–hole pair creation energy which means less carriers induced by the same amount of energy transferred by a radiation. For instance, in SiO_2 the electron mobility is about 20 $cm^2/V/s$ (versus 1400 $cm^2/V/s$ in Si) and $E_{eh} = 18$ eV (against 3.6 eV in Si). Consequently, *for SEE considerations* in which we assess the effects of injecting parasitic currents on critical circuit nodes, oxide charges are not the primary concern. If we were to focus on TID on the other hand, it would be quite the opposite: charges deposited inside active silicon, whether or not they trigger soft errors, eventually flow away or recombine and there remains to study the long-term influence of charge trapping in the insulators (or at the interfaces).

1.2.3 Circuit response – classification of soft errors

Following up on our bottom-up description stated in Table 1.1 on page 11, we can now discuss the phenomenology of SEEs at circuit level: once the deposited charges transport in silicon and are collected at circuit nodes, the circuit responds to this excitation and soft or hard errors may occur. This can be modeled by introducing radiation-induced stimuli in analog circuit solvers like the universally-famous Simulation Program with Integrated Circuit Emphasis (SPICE) from Berkeley, and its numerous derivatives such

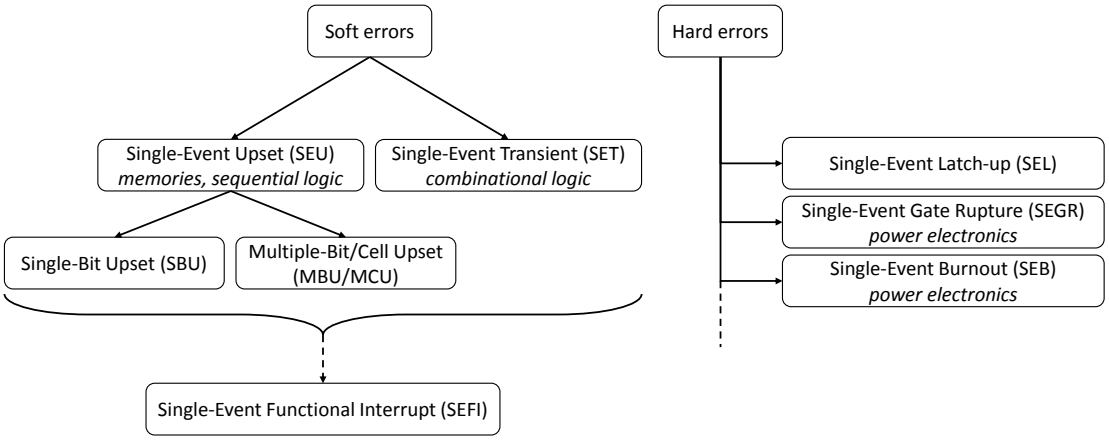

Figure 1.13: Classification of Single-Event Effects (SEE). Adapted from [36].

as HSPICE from Synopsys [35] or Eldo by Mentor Graphics [8]. As mentioned earlier, the circuit's feedback can have a strong influence on radiation-induced charge transport, which itself modifies the electrical state. This is because the time constants mentioned above for drift and diffusion currents are comparable to characteristic circuit response times (e.g. the standard inverter delay is only a few dozen picoseconds in modern technologies). In Chapter 3 we will shed more light on the implications of this, both in terms of phenomenology and modeling challenges.

Figure 1.13 defines some appropriate terminology for SEEs. Among hard errors, Single-Event Gate Rupture (SEGR) and Single-Event Burnout (SEB) are mainly related to power devices with high-voltage or high-current constraints (the former designates the breakdown of a transistor's gate dielectric from a radiation event, and the latter is defined as avalanche breakdown of a device when a radiation strike induces a high forward-bias state). However Single-Event Latch-up (SEL) can happen in usual CMOS circuits, when the parasitic n-p-n-p thyristor structure shown in Figure 1.14 is triggered by a radiation event: when minority carriers get collected at the junctions, majority carriers stay behind, which alters the well potentials. This can lead to a situation where the collector of the Bipolar Junction Transistor (BJT) formed by the pMOS drain/n-well/p-substrate structure feeds the base of the nMOS drain/p-substrate/n-well BJT, and *vice versa*. If the bipolar gain of these coupled parasitic BJTs is sufficient, this creates a self-sustaining low-resistance path from VDD to GND, i.e. a short circuit between the supply rails. Provided abnormally-high currents can be detected and avoided by powering down the chip, SEL may not be destructive—nevertheless, it is categorized as a hard error. Note that the *majority carrier currents*, and the subsequent *bipolar amplification* they can lead to, will be discussed in much more detail in Chapter 3. For the sake of brevity, we did not mention them in Section 1.2.2 discussing basic collection mechanisms in bulk silicon technology. In SOI technology though, they play a crucial role in the radiation response of elementary devices.

Now as far as soft errors are concerned, for digital CMOS logic which encodes information in voltages, the common question will be whether or not radiation currents induce a voltage drop that alters a signal significantly enough to cause a circuit malfunction. In other terms—at the risk of stating the obvious—harmful impacts are those which fulfill two criteria: they not only inject significant charges, but they do it on a sensitive circuit node. For instance, the power supplies may collect massive radiation

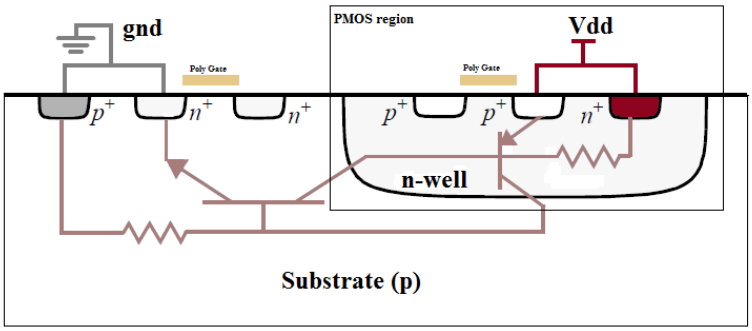

Figure 1.14: Single-Event Latch-up (SEL) occurs when a radiation event triggers the parasitic Silicon-Controlled Rectifier (SCR) formed by the n- and p-wells in CMOS technology. It is more likely under high voltage and temperature conditions. Taken from [37] .

currents, but they are able to sink an "infinite" amount of charge and oftentimes, said currents will not perturb the circuit. On the contrary, when a storage node of a bit-cell or sequential element (e.g. a flip-flop) changes electrical state due to a particle strike, a Single-Event Upset (SEU) occurs; when a combinational logic gate's output is disturbed by a radiation event momentarily, a Single-Event Transient (SET) can propagate in the fan-out cone of the gate; and when an SEU or an SET occurs on an architecturally-critical part of the circuit (for example a configuration register), it may translate into a Single-Event Functional Interrupt (SEFI). In other words, SEFIs are the manifestation of bad luck, where a particle striking at the wrong place and time triggers a large system malfunction. A SEFI can be undetectable: if a Phase-Locked Loop (PLL) configuration code is altered, dividing output frequency by two, the PLL may remain in this state without the circuit knowing that all operations are being performed twice as slowly.

Single-event upsets in memory cells and sequential logic

The most common kind of soft errors—and historically the first to have been discovered and characterized experimentally—are SEUs in memory cells and flip-flops. They were postulated in 1962 by Wallmark and Marcus [38], and as far as space electronics go, first observations were published in a 1975 issue of the Transactions on Nuclear Science, by Binder, Smith, and Holman who reported JK flip-flop upsets in a Hughes satellite [39].

Consider the typical six-transistor design of an SRAM cell shown in Figure 1.15a: the two central nMOS–pMOS pairs are cross-coupled inverters that form a bistable element, whose electrical state can either be high or low statically—e.g. a logical '1' at the input of the left inverter gives a '0' at its output, and the second inverter's feedback reinforces the '1'. The other two nFETs are access transistors enabling read and write operations. When the bit-cell is in "hold" (storage, retention) mode, the access transistors are turned off: the word lines are at GND. The bit lines are not loaded (both at GND), while during read or write operations they are complementary.

Given what we discussed in Section 1.2.2, we can analyze the radiation sensitivity of the bit-cell in hold mode. In bulk technology, the most radiation-sensitive areas arise in transistors with a finite $|V_{\mathrm{db}}|$ or $|V_{\mathrm{sb}}|$, so in general they will be located at the drain junctions of off-state MOSFETs [11] (the sensitive areas in SOI technology will be discussed in Chapter 3). For instance if BLTI is at VDD, the main charge-collecting p-n junctions will be the drains of PU2, PD1 and the BLTI electrode of PG1. Therefore

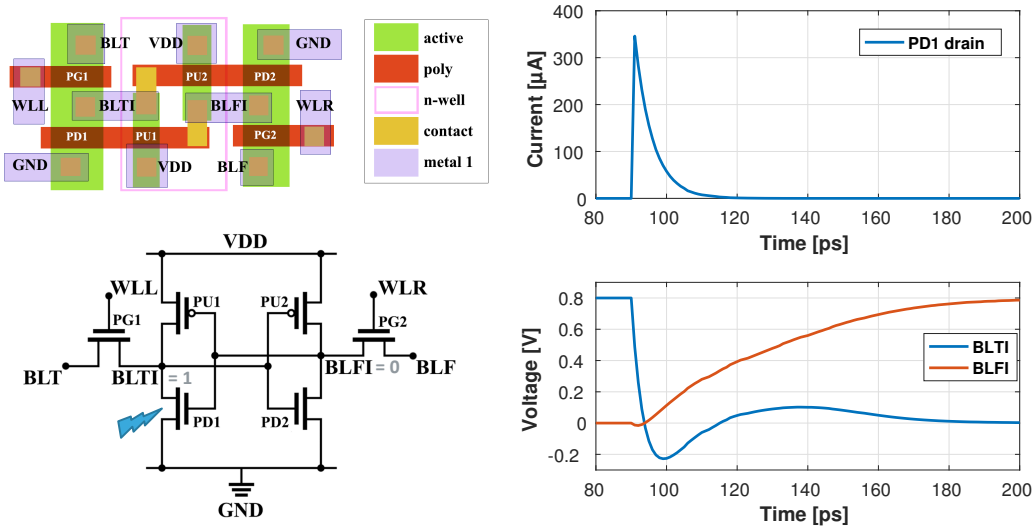

(a) (b)

Figure 1.15: SEU in a six-transistor SRAM following an nMOS impact: (a) Cell layout and schematic. "WL" and "BL" stand for Bit- and Word-Line. The 'I', 'T', 'F', 'L', and 'R' suffixes stand for Internal, True, False, Left, and Right respectively. "PU", "PD", "PG" are Pull-Up, Pull-Down, and Pass-Gate transistors – (b) Waveforms of the nMOS drain current and voltages inside the bistable element.

if a particle impacts nMOS PD1 as shown in the illustration, electrons are collected and BLTI starts to drop—conversely, if we were to strike pMOS PU2, holes would be collected and BLFI would increase. An SEU occurs if the voltage perturbation at the output of the struck inverter is significant enough (in amplitude and duration) that the following inverter sees a modified input that ripples to its output: as shown in the figure, BLTI and BLFI switch states and the bit of information is lost.

An important quantity to characterize radiation robustness is the notion of *critical charge*, defined as the integral of the smallest radiation current able to trigger an SEU:

$$Q_{\text{crit}} = \min_{\text{all upsets}} \int_0^\infty I_{\text{rad}}(t)\mathrm{d}t \tag{1.9}$$

The critical charge is a *circuit property*, specific to each node that can be permanently flipped, and depends on which transistor is hit. For the symmetrical SRAM considered here, the critical charges of BLTI and BLFI are equal. Nevertheless, the circuit response will be different for an nMOS impact lowering BLTI or a pMOS impact raising BLFI. For such small circuits as bit-cells though, "the" critical charge is usually understood as the smallest value among all scenarios, thus quantifying the most sensitive case. In today's technologies, SRAM critical charges lie in the femtocoulomb range, i.e. to trigger a soft error, only a few thousand electrons can suffice. Roughly speaking, it scales as:

$$Q_{\text{crit}} \approx C_{\text{node}} \cdot V_{\text{dd}} + I_{\text{restore}} \cdot \tau_{\text{flip}} \tag{1.10}$$

because as analyzed in [40], the storage capacitance needs to be (dis)charged and the restoring device has to be overpowered for a sufficient duration. For instance when PD1 is struck and BLTI drops, 'on' state PU1 starts driving current to raise BLTI back up with a restoring force $I_{\text{restore}} \sim I_{\text{on}} = I_{\text{d, sat}}$ close to the saturation-regime current ($|V_{\text{gb}}|$

and $|V_{\text{ds}}|$ are both "large" during that period of time).

In a nutshell, a good part of the soft-error paradigm amounts to the notions of collected and critical charges:

- Assessing the robustness of a given circuit, whether experimentally or by means of simulations, essentially means figuring out how Q_{coll} and Q_{crit} compare to one another in a particular radiation environment and circuit configuration. At this point, note that since Q_{crit} decreases linearly with voltage according to (1.10) while Q_{coll} decreases less-than-linearly (see (1.7) on page 18 and subsequent discussion), circuits are typically *more sensitive to radiation at low voltage*.
- Engineering *radiation-hardened* or rad-hard circuits is the task of making Q_{coll} lower and/or Q_{crit} higher. As said previously, Q_{crit} is a circuit property, thus it can be increased with schematic-only solutions (e.g. with larger node capacitance, regardless of the actual physical implementation). On the other hand, reducing Q_{coll} is necessarily a matter of lower-level physics; techniques which act upon the circuit's layout to reduce charge collection also belong to Radiation Hardening By Design (RHBD) techniques. For instance, placing sensitive transistors away from each other leaves the schematic unmodified but may help mitigate undesired *charge sharing* effects. Finally, if the technology process itself is altered to achieve robustness, then the method is called *hardening by process*. For instance, making use of heavily-doped layers may favor recombination of unwanted radiation-induced carriers. However in bipolar devices for instance, short carrier lifetimes also mean lower gain, thus trade-offs are necessary to limit the impact on device performance. In short, *robustness always comes at a cost*.

Our discussion of SEUs has been SRAM-centric for exposition purpose, but note that in fact, SEUs can occur in many kinds of sequential logic elements such as D-type Flip-Flops (DFFs): unless for read-only cells, a "true" memorizing function is usually achieved by a feedback loop—unlike for instance in DRAM where the storage capacitors eventually leak out the information in parasitic resistors. Such feedback loops can always be overpowered given a sufficient amount of charge, which roughly speaking corresponds to the write-cycle energy.

Single-event transients in combinational logic gates

Now, another kind of soft error that may arise in a logic circuit is an SET, or a temporary glitch at the output of a *combinational* logic gate: when the transistors are wired in a purely *feed-forward* fashion, if a node potential drops following an ionizing impact, the circuit is bound to recover its initial state at some point—the cumulative effects of TID or DD will only be felt after many more particle hits at the "same" place. Issues may arise, however, if in the meantime the voltage pulse propagates to downstream elements and for example, if this incorrect data is sampled by a flip-flop. In other words, for a digital SET to trigger visible effects, the transient has to reach the end of a data path on the rising edge of the clock signal. Such a scenario may seem quite unlikely and in fact it is, which is precisely why single-event upsets were observed prior to the discovery of single-event transients: SETs were theorized upon in 1983 [41], and were observed two years later [42]. Several kinds of *masking* can cause a radiation-induced voltage transient in a logic path to be overall "silent":

- Logical masking (Figure 1.16a) happens when the propagating glitch is made silent

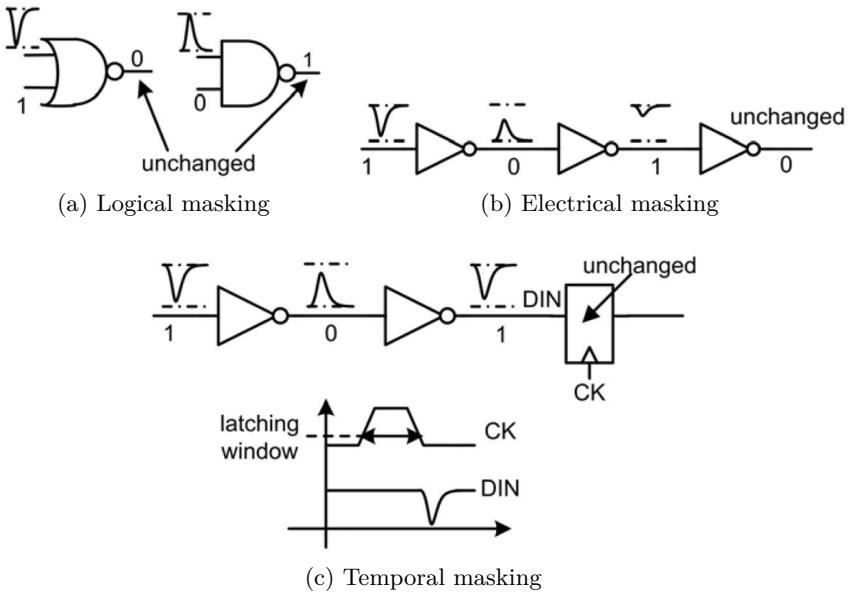

(a) Logical masking (b) Electrical masking

(c) Temporal masking

Figure 1.16: SET masking phenomena in combinational logic. Taken from [43].

by the boolean equation of a traversed logic gate. For instance, a NOR gate whose first input is high will block any negative SET on the other pin.

- Electrical masking (Figure 1.16b) occurs when a short transient—i.e., with large bandwidth—is filtered by comparatively slower logic stages—i.e., with lower cut-off frequency. In that case, along the logic path the glitch amplitude will attenuate and its width will shrink, until it eventually vanishes. Note that while some logic gates can narrow SETs, others may also exhibit SET broadening, depending on their rise and fall times balance. Also, logic gates being non-linear voltage amplifiers with significant gain, a glitch with sufficient initial swing (typically $V_{\mathrm{dd}}/2$ for balanced gates) will be amplified, or in other words rectified into a true digital signal.
- Temporal masking (Figure 1.16c), as explained above, is the fact that a propagating pulse will not lead to a soft error if it reaches the end of a combinational path outside the latching window of the sequential element ending the data path.

These different masking factors can strongly depend on the Process, Voltage, Temperature (PVT) operating point, and furthermore on the manufacturing process; in Section 1.5.1 we will discuss how SET masking is affected by technology downscaling.

1.3 Characterizing and quantifying single-event effects

Having presented the chain of physical processes leading to single-event effects, we now discuss how circuits are experimentally measured for soft errors, and the metrics that can be extracted to characterize the radiation sensitivity of an IC.

1.3.1 Accelerated testing in particle beams

Soft errors are random phenomena by nature. To qualify their occurrence frequency, or the Soft-Error Rate (SER) of a given circuit for a chosen mission, the most faithful

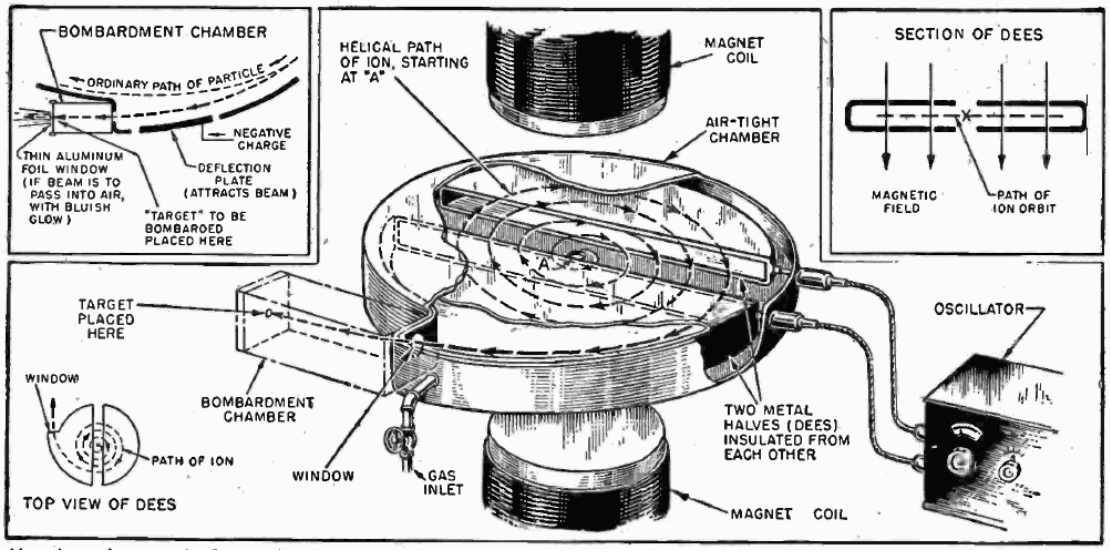

Figure 1.17: Diagram of a cyclotron, taken from [45]. Size of the magnets has been kept down to show dees path of the ion.

experiment would be to test it in its intended environment and with stimuli that are most representative of the mission profile. However, such stress conditions are quite impractical: on Earth, soft errors are rare, thus real-time testing experiments need to be performed on many replicas of the same circuit in order to collect meaningful statistics and derive an SER: in [44] for instance, 7 Gbit of SRAM had to be monitored for over a year to record about one hundred events. For space-bound electronics, the SER may be larger due to much-increased particle fluxes, but the associated volumes of chips are much lower than for mass-market applications. Besides, due to the cost, rarity, and non-zero risk of rocket launches, in most cases flying test chips to orbit beforehand is not an option, and only the finished product sees the real environment.

As a result, for practical assessment of the SEE sensitivity of circuits, one must resort to particle accelerators whose fluxes are orders of magnitude above natural ones[5]. The most common type of accelerator used is based on the concept of cyclotron, invented by Ernest O. Lawrence in 1932 (Figure 1.17): charged particles placed in the uniform, constant field of an electromagnet follow a circular trajectory whose period does not depend on the radius—Lorentz force $Qv \times B$ (with conventional notation) can only bend a particle's motion but not increase its velocity or energy. Then by tuning the pulsation of an alternating voltage applied to electrodes lying in the plane of trajectories (the "dees", after their shape), a synchronous RadioFrequency (RF) electric field arises, causing particles to accelerate in outward spirals. They finally reach the output beam at radius R with kinetic energy $Q^2B^2R^2/2M$ (again with usual notation). For ions at a given mass-to-charge ratio[6], the resulting energy per nucleon will thus be constant. The RADiation Effects Facility (RADEF) in Finland [46, 47], for instance, delivers a heavy-ion cocktail of 9.3 mega-electronvolt per atomic mass unit (amu) with species from

[5]Highly radioactive sources can also provide high acceleration factors compared to natural environments, and for instance thorium or americium sources with kilobecquerel or megabecquerel activities are used in alpha-particle testings for terrestrial applications. However, the modest energies emitted by such radionuclides (in the mega-electronvolt range) are not sufficient to emulate cosmic rays and qualify chips for space compliance.

[6]Particles with identical mass-to-charge ratio follow the same trajectory under any electromagnetic field, so for instance a deflecting magnet can select ions according to M/Q, but not M or Q separately.

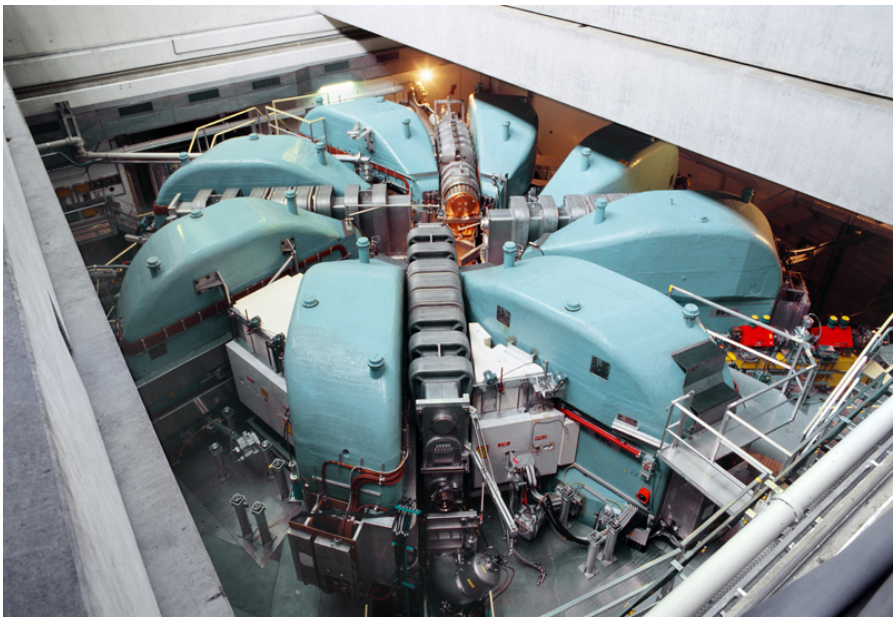

Figure 1.18: Proton accelerator at the Paul Scherrer Institute. Photograph: Markus Fischer.

nitrogen (Z=7, A=14) to xenon (Z=54, A=131). This allows to cover an LET range from 1.83 to 60.0 MeV$\cdot$cm^2/mg in silicon, with typical fluxes around 10^4 ions/cm^2/s to reach fluences of $10^6 - 10^7$ ions/cm^2 for a run lasting a few minutes. Throughout this thesis, we will also frequently discuss measurement results obtained with high-energy protons at the Paul Scherrer Institute (PSI) in Switzerland [48, 49], whose cyclotron is shown in Figure 1.18 and can typically deliver 10^{11} protons/cm^2 within half an hour. These beamline facilities, among others, were audited and approved for compliance to certain standards, for instance to qualify for radiation testing of space-bound electronics. In Europe, the standard in use is ESA/ESCC basic specification No 25100 [50], which also formalizes recommendations for measurement protocols and metrics to normalize single-event characterizations. An important constraint, for example, is the obligation to monitor the total dose even during single-event testing, as SEEs can be affected by TID. For the figures given above, doses can be as high as a dozen kilorad or a few hundred gray (a gray is one hundred rad, or one joule per kilogram), which for the devices discussed throughout this text is not necessarily negligible.

Beside beam control apparatus provided by the accelerator facility, and the Device Under Test (DUT) itself, typical test setups include a test Printed Circuit Board (PCB) to mount the chip(s) on, regulated power supplies and connectors, volt- and ammeters, and a PC for overall control and monitoring during the runs. FPGAs and oscilloscopes are very common as well, allowing to perform customized and reconfigurable functions depending on the devices being irradiated; on the other hand, test board developments often have to be dedicated to one family of chips. Even when no specific test board has to be designed, several steps need to be taken before chips can undergo a radiation campaign: depending on the range of particles that will be used, packages may need to be opened, and parts may have to be further thinned (or backlapped) in case back-side irradiation is planned. Indeed, except for long-penetrating (atmospheric) neutrons or (space) protons, worst-case cosmic ray energies are not achievable within particle accelerators, and thus GCRs are only *emulated* based on their LETs but with shorter ranges—we further discuss this approximation in Appendix B. Then the Design of Experiments (DoE) needs to be planned, and considering the high cost of beam time, dry

runs can be carried out to stimulate the test bench before the actual beam slot. Thus, between "silicon out" and a radiation campaign, there is always an incompressible amount of time of several months. Together with the logistics of travel to the facility and the efforts associated with data post-processing and report authoring tasks, this means that beam experiments are expensive activities, with long timescales involved.

1.3.2 Test structures for single-event upsets and transients

In the frame of reference of a silicon foundry—i.e. the work environment in which this PhD program was set, radiation testings are usually performed on the essential building blocks of the technology under qualification. Characterizations at end-product level, on the other hand, can be very challenging in terms of observability. Tests on complex systems, when technically feasible, must be carried out jointly with the customer embedding the chips in a larger architecture. Therefore, to qualify the basic elements of a technology, inevitably memory cells will be irradiated, and to complete the SEU picture, sequential logic cells need to see some beam time as well. Provided SETs can also be characterized in combinational logic, a *representative set* of the digital blocks is described, allowing to derive SER estimates based on the content of the System on Chip (SoC). Note that the same goes for analog blocks such as phase-locked loops or power amplifiers, where oftentimes a single flagship Intellectual Property (IP) is qualified, allowing to mitigate the risks on neighboring IPs by inheritance.

Starting by order of simplicity, memory cells are easily amenable to radiation testing: based on a single address, it is possible to access any location in the array, and thus keep the observability on SEUs occurring in the irradiated bit-cells. Thus by creating large memory cuts (usually a few megabits for SRAMs), one leverages a high number of bit-cell instances which, combined with a high particle flux, can provide enough data for accurate characterization. For sequential logic, specific structures have to be developed in order to test many cells simultaneously. For flip-flops for instance, generally one will connect them in a chain to form a shift register whose clock signal is distributed to all cells. By inserting a given pattern at the chain input and verifying its coherence once it exits the chain, one can typically monitor dozens to hundreds of kilobits of sequential logic for SEUs. The test protocol can be more challenging than for SRAM cells however; static operation is possible (insert the pattern, expose, shift the sequence out, and compare), but oftentimes, at-speed dynamic testing is preferred since it creates more realistic stimuli. The reference and output patterns then need to be compared on the fly, via an external FPGA (which limits the operating frequency to a few hundred megahertz) or with on-chip logic. Even in the latter scenario, running above 1 GHz is difficult due to the current drawn by such test structures.

Now, unlike persistent bit flips, capturing evanescent pulses is a much more challenging task. Over the course of the 2000s, the research focusing on measurement of single-event transients gave rise to several publications. Direct measurement of SETs has been demonstrated, by wiring sensitive devices to RF Input–Output (IO) pads and using a high-bandwidth oscilloscope for acquisition [51]. In order to synchronize the trigger signal of the oscilloscope with the particle impacts, the method requires the use of an ion microbeam that allows to target a submicron focused spot with just a few particles, down to a single ion. Alternatively, the authors also performed pulsed laser irradiation, which poses the additional theoretical challenge of the equivalence between the photocurrent and the ion-induced current: the initial electron–hole pair distributions need to be matched, both in terms of spatial and temporal structures, to yield quantita-

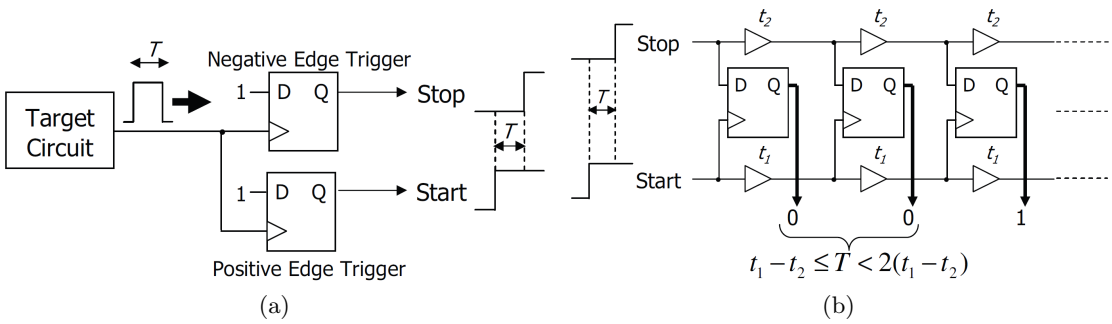

Figure 1.19: Vernier Delay Line (VDL) structure for transient width measurement: (a) Generation of edges at the start and end of the glitch – (b) Delay chains and flip-flops measuring the time difference between both edges. Taken from [55].

tive comparison between laser and ion irradiation results (see [52] for an excellent survey paper of laser testing for SEEs). Under a broad particle beam in regular accelerators, several test structures have been developed to measure the frequency of transient pulses, and their characteristics. When talking about digital SETs, apart from polarity (negative for a high–low–high or positive for a low–high–low glitch), the defining parameter to be measured is the pulse width. On-chip methods that were first developed allowed to quantify the SET duration by propagating it in a delay chain, and latching it along the way [53, 54]. The measured width is then a multiple of the delay element, imposing a physical barrier on the resolution at a certain process node, given by the standard inverter delay with Fan-Out (FO) of one. In 2010 however, it was shown that by using a Vernier Delay Line (VDL), sub-FO1 resolutions on the pulse width were achievable [55]: two chains of slightly different delay elements are used to propagate the SET edges. One chain distributes the clock signal to a set of DFFs, and the other is connected to the data pin of the flip-flops. Thus, the SET edges will race in both delay lines, and trigger a sequence of bit flips in the DFFs whose length is equal to the pulse width divided by the delay *offset*. Owing to this improved resolution, VDL-based solutions are now very popular for on-chip pulse width measurement.

1.3.3 Metrics for single-event effects

When irradiating the test structures presented above, due to the random occurrence of SEEs, the data acquired allows to derive *statistical aggregates* that quantify the radiation response of the technology. The quantity of choice, of course, is the SER, but for space electronics, direct measurement of the natural error rate is difficult. Firstly, real-time testing is impractical before satellite launch, as mentioned previously. Secondly, for terrestrial applications, accelerated testing is enabled by broad-spectrum sources reproducing the energy shape of a normalized environment: a representative NYC spectrum was standardized by Joint Electron Device Engineering Council (JEDEC), and is well matched by sources at Los Alamos Neutron Science CEnter (LANSCE) [56] or TRI-University Meson Facility (TRIUMF) [57] for example. For space missions on the other hand, the radiation environment can greatly vary from one satellite trajectory to another, and depending on space weather as well. In other words, the variety of possible space environments is such that accelerated testing cannot directly yield the SER for a given orbit, but rather serve to quantify the circuit's intrinsic sensitivity under a mono-energetic, single particle type, unidirectional beam, over a range of energies, nuclei, and

directions. Then knowing the spectral content and direction dependence of the targeted environment, this multivariate sensitivity can be reweighted so as to project an SER.

This is formalized in the notion of *single-event cross section*, defined as:

$$\sigma = \frac{N_{\text{evt}}}{\Phi} \tag{1.11}$$

where N_{evt} is the number of recorded events, and fluence Φ is the number of particles divided by the irradiation area:

$$\Phi = \frac{N_{\text{part}}}{A_{\text{irrad}}} \tag{1.12}$$

By combining (1.11) and (1.12), it is easy to see that the cross section measures the *fraction* of particles able to trigger an event. These definitions are completely identical to those established by the ICRU for radiation–matter interactions such as nuclear fissions, photon absorptions, elastic electron scattering events, and so forth [20]; essentially, the single-event cross section of a chip simply means we are shifting the notion of "interaction" towards a circuit-oriented perspective. It thus represents the sensitive area that the chip effectively opposes perpendicularly to the incoming flux.

Then provided no dose effects intervene, we can differentiate (1.11) to find that:

$$SER = \frac{\partial N_{\text{evt}}}{\partial t} = \sigma \cdot \frac{\partial \Phi}{\partial t} = \sigma \cdot \phi$$

where ϕ is the flux and SER is expressed in events per unit time. Thus, under a constant flux neither the cross section nor the error rate vary with time, and a plot of N_{evt} versus time will be linear, when seen from a distance to overlook statistical fluctuations. Now, for rigorous definitions, as mentioned, the cross section should only be considered at a given energy, for a unique ion species, and at fixed angles, although we did not explicitly write these dependences in (1.11). Thus, in the general case, the error rate can be retrieved by folding the multivariate cross section with the environment flux, cast in the same variables:

$$SER = \sigma * \phi \tag{1.13}$$

where operator $*$ stands for convolution. By "same variables", we mean that under certain hypotheses, the dimensionality of this convolution product can be reduced: assuming isotropy, the angular dependence vanishes; neglecting detailed ion track structure, the Z summation (for atomic number) and E integration (for energy) hidden in (1.13) can be replaced by an integral over the LET. All these will be discussed in Chapter 4 when we address on-orbit upset rate calculation (Section 4.3 page 121).

To summarize, the number of measured (or simulated) events provides direct insight as to the radiation sensitivity of the tested circuit, whether it be via the cross section or the SER. The occurrence of SEEs is nondeterministic, and on a mathematical note, this is usually modeled as a Poisson process—the least-assumption model for a stochastic counting process, leading to confidence intervals given by a chi-squared distribution (χ^2). Dividing (1.11) by the number of measured instances, one can also derive a sensitivity per bit, per logic cell, per IP, or even per chip for large systems. Thus, there are three possible levers for ensuring that the circuit is qualified with good statistical confidence, as suggested in dotted lines throughout the two previous sections:

- First, to increase the flux, at the risk of magnifying non-realistic effects such as multi-impact events, with too close correlation in space and/or time: when using dramatic

acceleration factors, the system may not recover quickly enough in-between two events (to only mention proximity in time), leading to pessimistic SER estimates.

- Then, to increase the exposure time, at the risk of inducing dose effects—and with practical limitations also since particle beams are charged by the hour.
- Thirdly, to increase the number of instances. This can be done on one test chip within the allotted silicon area, but with more complicated IC design: for example and as mentioned earlier, the IR drop induced by power consumption from large shift registers can be challenging. Alternatively, multiple DUTs may be used (which erases dosing risks), but with the difficulties that it entails in terms of test protocols. For instance, test boards without a DUT socket require the chip to be soldered on, thus swapping DUTs more than a few times is not always an option.

1.4 State of the art in soft-error simulation

The previous section has been devoted to experimental characterization of single-event effects. Now, in the field of radiation effects just like in any other scientific domain, a large body of work is dedicated to the development of accurate models, and their implementation and use in simulators. In an attempt at categorizing applications that are clearly intertwined, we may consider that in applied fields the need for modeling and simulation is impelled by three main axes: to understand, predict, and optimize.

Indeed, in a simulation, virtually any quantity of interest can be probed, while in the particular field of radiation effects, the observability during measurements is often hindered by the particle accelerator facility and its stringent environment. Physical insight can thus be gained from accessing fundamental variables such as those governing charge transport at nanometer scales, thereby providing an intuitive compass for reasoning. Then provided the semiconductor modeling is close enough to the real device behavior, a simulator also equips us with large prediction abilities: just knowing the equations is not enough, and computational physics yields intelligibility to complex processes where a human's best guess would not be tangible. From thought experiments and "what-if" simulations, we can project the behavior of devices that have not been measured, or have not yet been manufactured altogether. Finally, with this ability to anticipate on the performance of untested circuits, time can be saved compared to design iterations with silicon validation at each step; therefore, more optimization loops can be performed so as to tailor a design to best suit its requirements.

The simulation tools developed and used in this work serve these three inclinations: in the next paragraph we present TCAD as our tool of choice for understanding the peculiar charge transport processes at work in our studied devices; then we introduce the general architecture and examples of state-of-the-art Monte Carlo radiation simulators used in the community for SER prediction and RHBD optimizations.

1.4.1 TCAD or detailed semiconductor simulation for transistor-level analyses

After the high-energy physics processes leading to the "initial" electron–hole pair distribution in silicon, semiconductor physics comes at play to describe the device response. The constituting equations of electromagnetism and carrier motion are tightly coupled, and hence the formulated problem is inherently stiff: as soon as charges start moving

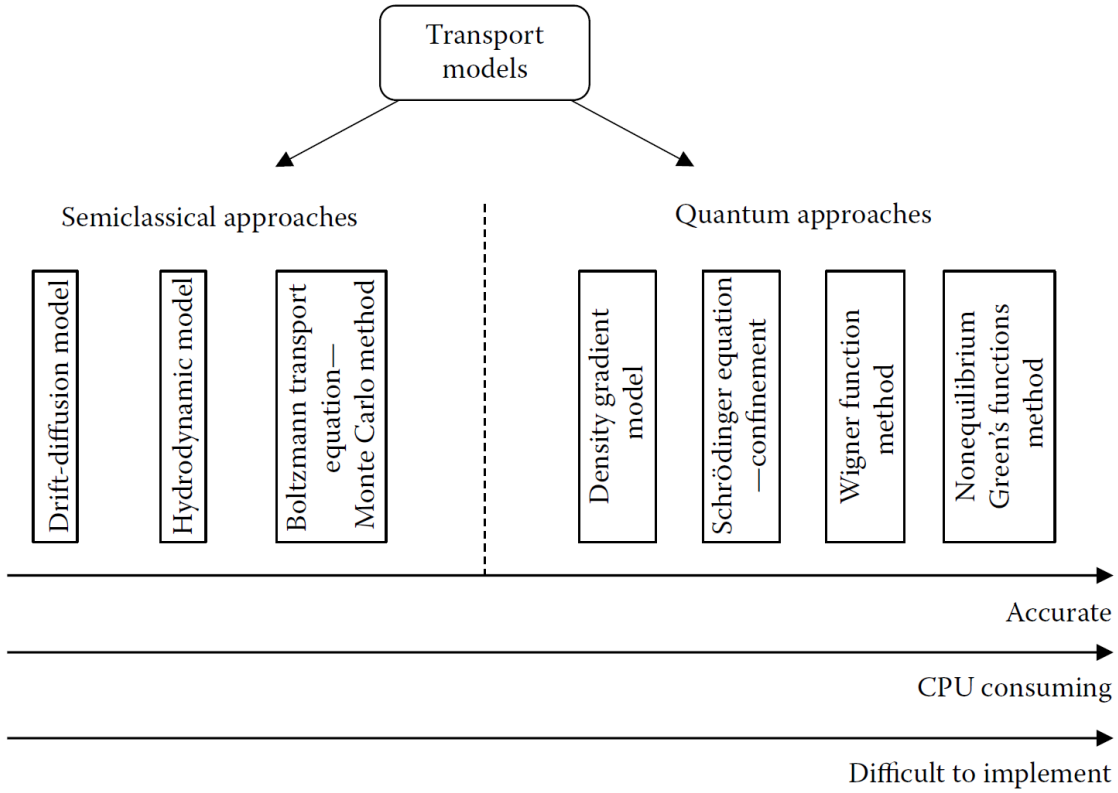

Figure 1.20: Carrier transport models for semiconductor simulation. Taken from [36].

in the electric field, they redistribute the field, and thus the system of equations for electrostatics and carrier transport must be solved self-consistently. Apart from RF or optical devices, the electromagnetic problem is always reduced to electrostatics described by Poisson's equation; carrier transport models, on the other hand, vary in speed and accuracy, and a panorama of them is given in Figure 1.20.

Aside from full quantum mechanical treatment that is only necessary for purely quantum devices such as Josephson junctions or nanodots, semiconductor models are typically based on the semi-classical Boltzmann Transport Equation (BTE). The BTE describes the density of carriers both in real space and reciprocal space, i.e. the space of wave vectors or k-space. It is semi-classical in two senses: it first inherits from quantum mechanics because the semiconductor band structure derives from Schrödinger equation, and its energy occupation statistics is dictated by Pauli's exclusion principle leading to a Fermi–Dirac distribution at equilibrium. Then, the BTE is a semi-classical formulation in the sense that carriers are treated as wave packets of well-defined position and momentum: they are close to saturating Heisenberg's inequality, i.e. extending over a few lattice cells whilst having small spread in the crystal band structure. The transport equation then describes the evolution of this carrier density under the scattering mechanisms relevant to condensed matter physics: electrons can bounce against ionized impurities depending on the doping level, they can undergo intervalley scattering at high electric fields, or they may scatter against phonons depending on the temperature—to just illustrate how PVT corners can influence charge transport. Practical solutions of the Poisson–Boltzmann system are often computed with the Monte Carlo method, whereby the trajectories of individual electrons are simulated as free flights (ballistic transport) under the local electric field, interrupted by the aforementioned scattering (or recombination) events. In the general case, one must thus synchronize all carriers after a certain

time step, and update the electric field for the next step. Such self-consistent "ensemble Monte Carlo" solvers are highly numerically-intensive, and for instance computing the current–voltage characteristic of a MOSFET may take several days.

Simplifications to the BTE are thus necessary for computationally-efficient simulation of microeletronic devices; the simplest model used in commercial TCAD tools is the well-known drift–diffusion approximation, which was first derived by Van Roosbroeck in 1950 [58]. The model can be obtained from thermodynamics considerations, or derived from the first-order moment of the BTE (i.e. multiplying it by carrier velocity v and integrating). Carrier transit is assumed to be slow before the energy relaxation time, and then the resulting current is the sum of two terms: first, a *drift current* parallel to carrier velocity $v = \mu \cdot \mathcal{F}$ imparted by the electric field $\mathcal{F}$ to the free charges of mobility μ; then, a *diffusion current* proportional to $D \cdot \nabla c$ where c is the carrier concentration (i.e. n or p for electrons or holes) and D is their diffusivity. The diffusion current is analogous to a fluid description and tends to even out the concentration gradients. Now, all these assumptions can break especially for short-channel devices where non-stationary effects can arise, such as velocity overshoot (above saturation velocity) or even ballistic transport for deca-nanometric CMOS devices. These can be tackled by the more elaborate hydrodynamic model, derived from the second-order moments of the BTE (i.e. this time multiplying by carrier energy or v^2 the velocity squared). Energy-balance equations are then obtained, meaning that the carriers' energy is no longer perfectly aligned with the electric field, and their temperature is not necessarily that of the crystal lattice. Both the drift–diffusion and hydrodynamic models are now commonly implemented in TCAD software, and provide solid foundation for transistor-level studies of charge transport. On the other hand, Monte Carlo Boltzmann solvers are typically implemented in more fundamental research tools, and used to deliver effective models for use in TCAD suites. For example, mechanical stress effects, impact ionization under high fields, or in fact all mechanisms where a full-band description is necessary, can be studied at length with BTE solvers; they are then accounted for in TCAD with an effective strained mobility model or a properly-parametrized avalanche generation rate.

The study of single-event effects in TCAD is but a particular case of transient semiconductor simulation: after the virtual device has been brought to a steady state by ramping its terminal voltages under quasi-stationary conditions, a time-varying solution is computed where charge carriers are injected via a radiation-induced carrier generation rate that best reproduces the ion track, and whose typical features were underscored in Section 1.2.1 on page 13. The space dependence of the electron–hole pair generation typically defines a straight track at a certain LET, with Gaussian transverse shape of a few dozen nanometers; the time dependence of the generation rate is usually Gaussian as well, with characteristic time in the picosecond range. To properly capture these small scales and quick variations, the simulation mesh needs to be refined locally both in space and time, and caution must be taken to ensure numerical convergence. Over the years, various milestones have been set on the scale of the achieved SEE simulation domain, with increased dimensionality (2D at first, then 3D) or transistor count. For instance, the first six-transistor SRAM cell modeled within a contiguous 3D domain, was only reported in 1998 [59]. The alternative to full TCAD simulation is to use *mixed-mode* abilities offered by most tools, where the device equations are coupled via the boundary conditions to circuit equations making use of compact models. In other words, the impacted device is simulated in the physical domain, while other devices are solved in the electrical domain with SPICE models. This greatly alleviates computational costs, while sacrificing some accuracy when charge sharing mechanisms are overlooked—in that case,

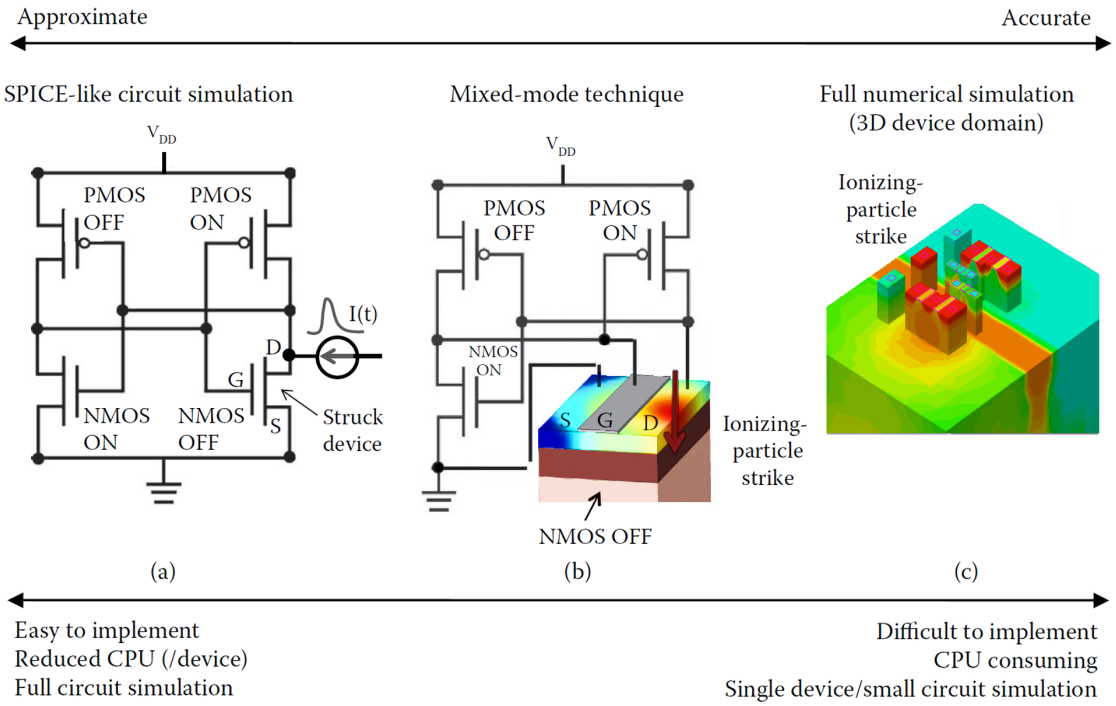

Figure 1.21: Abstraction levels for SEE simulation, from circuit to device domain [36].

including neighboring devices in the physical domain is necessary. The two simulation approaches are illustrated in Figure 1.21b and Figure 1.21c. Overall, typical TCAD simulation times range from a few minutes to a few hours for one ion impact, depending on the number of mesh vertices which itself depends on dimensionality, device count, domain size, and mesh refinement. A vast corpus of simulation-centric studies has been published and will be elaborated on throughout this text; if only to cite one prominent work, [11] can be referred to as a remarkably exhaustive study of SEU mechanisms in both bulk Si and SOI technologies.

To conclude this overview of SEE simulation with TCAD, note that although these models are based on fundamental semiconductor principles, they are not *ab initio* approaches: as suggested above, full-band Boltzmann solvers can be leveraged to create effective models to fix certain coefficients in TCAD transport models; but even then, there still exist some free parameters. Typically, this means that TCAD simulation decks have to be calibrated against electrical characterization on existing devices, before they can be used for explorations on virtual devices that do not extrapolate the simulation too far outside its validity range. For MOSFETs for instance, the electrostatics can be validated based on the $C_{gg} - V_g$ characteristic, and transport is calibrated from $I_d - V_g$ curves. Gate stack parameters allow to fine-tune the electrostatics (dielectric permittivity or metal work function), and interface properties (e.g. trap densities affecting the inversion-layer mobility) can adjust carrier transport.

1.4.2 Monte Carlo error rate prediction for circuit analyses

After investigating the radiation response of elementary transistors in TCAD, oftentimes the goal will be to assess the overall sensitivity of a particular circuit, and possibly optimize its rad-hard design. Benchmarking is then based on the statistical metrics given in Section 1.3.3, page 30, namely the circuit cross section or error rate in the

targeted environment. Predicting these aggregates by simulation thus implies the ability to compute the circuit response to a large number of impacts, which imposes stringent constraints on the speed of SEE models: if one wishes to integrate over all dimensions of "particle space" by simulating several thousands of impacts, spending no more than a few seconds computing one ion strike is usually the requirement. This constitutes one of the major challenges addressed during this PhD, both from a scientific standpoint—i.e. having to find simplifying assumptions at minimal accuracy lost—and from a technical perspective—i.e. in terms of efficient implementation. Direct SER estimation with general-purpose TCAD software is, although not entirely intractable, incredibly CPU intensive. It was demonstrated in [11] for instance, but mainly as a one-off detailed study rather than a standard simulation procedure.

Because of that, the first techniques for SER projection were derived from drastic assumptions to offer mainly-analytical models. The historical Rectangular ParallelePiped (RPP) model [60], which was introduced in the early 1990s, considered constant-LET cosmic rays impinging on a box-shaped sensitive volume defined by a certain critical charge Q_{crit} for SEU threshold. Comparing Q_{coll} in (1.7), page 18, to this Q_{crit} and integrating over all directions and LETs present in the environment, it was then possible to derive an analytical SER based on a precomputed distribution of chords lengths in the RPP. This is obviously very simplified, both with respect to energetic particle transport (e.g. nuclear interactions are overlooked) and semiconductor physics (e.g. dynamic circuit properties cannot be captured with a static Q_{crit}). Since then, many research groups have been pushing the envelope to relax as many assumptions as possible, and reach the envisioned goal of coupling together state-of-the-art tools for each of the physical processes involved in complex SEE phenomena.

Roughly speaking, computational SER prediction (as opposed to mainly-analytical models) has been a research topic from the 2000s onward. Most simulators use Monte Carlo integration to explore the space of parameters defining particle events, e.g. not by drawing impact locations on a regular grid, but randomly instead to avoid complex imbricated loops. The chain of physical processes exposed throughout Section 1.2 is then simulated, at varying degrees of accuracy and speed. For instance, successive developments in the Monte Carlo Radiative Energy Deposition (MRED) code, developed at Vanderbilt University and NASA/GSFC, provide a fair account of the technical challenges that have been overcome throughout the years, as illustrated in Figure 1.22. The first work that used an early version of the code was published in 2003 [3] and demonstrated the use of Geant4 for energetic particle transport. Assuming an RPP and Q_{crit} criterion for SEU decision, this allowed to capture rare events thanks to detailed ion track simulation (i.e. down to individual δ-rays as opposed to collapsing the track structure to a single LET parameter). In 2009 [61], a large flip-flop structure with multiple sensitive volumes was simulated, this time linked to a SPICE solver to determine the occurrence of upsets, as shown in Figure 1.21a. These two works demonstrate how cascaded mechanisms of high-energy and semiconductor physics are tackled with increasingly accurate computations. Nowadays and thanks to its widely-adopted online interface CRÈME-MC (where MC stands for Monte Carlo), MRED has become the *de facto* standard for many fundamental SEE investigations.

Another contribution that ought to be mentioned, is MC-ORACLE developed at Université Montpellier-2 [4,62]. The "Diffusion–Collection" charge transport model used in MC-ORACLE originated in the early 2000s in [7], and solves the ambipolar diffusion equation in neutral regions via an analytical Green function approach, to which follows collection by assuming a certain drift velocity at the drain contacts. Since then, the

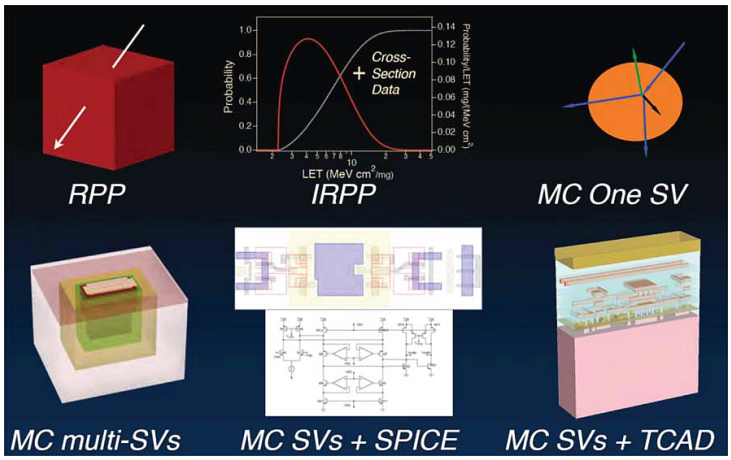

Figure 1.22: Hierarchy of approximations for single-event effects in MRED, from Rectangular ParallelePiped (RPP) to full Monte Carlo (MC) simulation using multiple sensitive volumes (SV) with TCAD in the loop. Taken from [65].

Diffusion–Collection model has been successfully ported in several Monte Carlo SER prediction platforms for planar bulk technology, for instance MUlti-SCAles Single Event Phenomena Predictive Platform (MUSCA SEP3) [63, 64] developed at ONERA, and including also the very tool discussed in this dissertation.

At this point, giving an exhaustive account of all existing simulators, along with their development milestones, would be irrelevant. Tackling this task methodically, is a large anthology article that was written in 2013 [65]. In the next chapter, we will describe at length the simulator developed at STMicroelectronics and IM2NP lab, that we use, maintain, and upgrade throughout this work. It is called Tool suIte for rAdiation Reliability Assessment (TIARA) [44, 66, 67], and after the detailed presentation of its code given in Chapter 2, in Section 2.6 we will be able to present a comprehensive comparison of TIARA with other state-of-the-art simulators. To conclude this presentation of Monte Carlo SER simulators, let us state once again that the purpose of such codes is to predict, and optimize: provided SEE models are fast enough compared to general-purpose TCAD, the ultimate goal is to obtain quick feedback during circuit design, allowing explorations for best-in-class hardening solutions. Another related application of numerically-efficient SEE models is the ability to quickly screen out sensitive nodes in large designs. Integration of SEE waveforms in an industrial CAD flow [68] can allow circuit designers to analyze Very Large Scale Integration (VLSI) circuits for weaknesses. This can be as important as SER prediction by the radiation expert on modest circuit sizes of typically a few standard cells.

1.5 Industrial context and research objectives

To complete this introductory chapter, the motivations for this work will now be highlighted. Recent CMOS technology trends and evolutions in the space market will be discussed, and after presenting the technology platforms that we focus on, the research objectives of this PhD will be stated.

1.5.1 Impact of technology node scaling on soft errors

The most prominent microelectronic technology trend that needs to be examined is of course CMOS technology downsizing. Based on the observation that radiation is by nature extrinsic while CMOS critical dimensions keep reducing—although at a lesser pace than the original Moore's law nowadays—with technology downscaling one should expect SEEs to obey certain rules.

Memories and sequential logic in planar bulk Si

Starting with SEU phenomena, historically the SER *per bit* has proved to decrease with integration [69]. This is the complex result of three main competing factors, namely the critical and collected charges, and overall sensitive area. Analytical derivations would be quite approximate here, but hints at the different trends can be given:

- Q_{crit} directly defines LET_{th} or the threshold LET, and as both transistor dimensions and supply voltages have been reduced, it has proved to drop steadily. In (1.10), page 24, and assuming constant-field scaling for instance, the critical charge scales quadratically with the process node.

- Q_{coll} (for a single cell) on the other hand, reduces more slowly with technology integration, as suggested by (1.7) (page 18): for an ideal p-n junction for example, the depletion width varies as the square root of voltage and substrate doping.

- The overall sensitive area, which can be defined as the maximum, or *saturation* cross section σ_{sat} at high ionizing power, scales quadratically upon dimension reductions.

How these factors combine into an upset rate can only be predicted via (1.13), page 31. Thus, what the observed SER trend shows is that, as integration goes, even though the collected charge exceeds the critical charge more easily, this increased sensitivity is outweighed by the reduced critical dimensions. Nevertheless, what [69] also shows, and as confirmed for instance by [70] in 65 nm bulk Si, is that Multiple-Cell Upsets (MCUs) can dominate the overall upset rate in deca-nanometric bulk technologies (Figure 1.23): the SER per bit may diminish, but a single ion impact may affect many cells via charge sharing effects. Thus as the packing density is increased, the SER *per chip*, if left unmitigated, can lead to unacceptable error rates even at sea level [71]. Consequently, on a critical memory matrix it may be necessary to use an Error Correction Code (ECC) to ensure acceptable error rates.

As for sequential logic cells, many RHBD solutions have been proposed. The Dual-Interlocked storage CEll (DICE) architecture, for instance, interleaves the storage nodes of latch elements in such a way that an SEU cannot occur if only one MOSFET is hit [72]. This is a very efficient SEU eliminator, but like any hardening strategy, it comes at a certain Power, Performance, Area (PPA) cost: the DICE approximately doubles the transistor count, and thus the cell area and power consumption. It has minor delay penalties however, since transitions on the redundant nodes occur simultaneously. In terms of scaling, similarly to MCU mechanisms, the DICE hardening scheme is challenged at advanced nodes: carriers can in fact be collected by more than one sensitive transistor, and the redundancy mechanism can be defeated [71]. Overall, in planar bulk technology, containing SEUs can require extra caution at highly-integrated nodes.

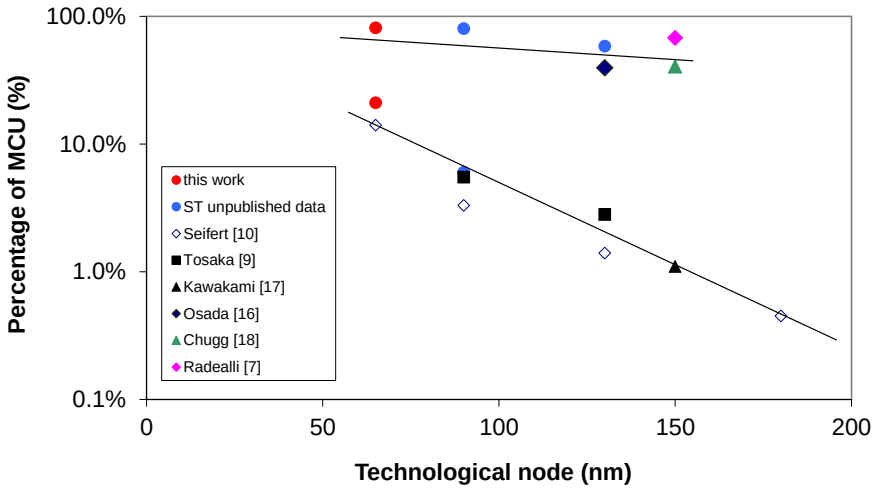

Figure 1.23: Multiple-Cell Upset (MCU) proportion increasing with density. Taken from [70].

Combinational logic trends

In combinational logic gates, the same considerations of $Q_{\text{crit}}/Q_{\text{coll}}/\sigma_{\text{sat}}$ apply for transient pulse generation. But for a complete analysis, one must also consider how SET masking factors scale across process nodes:

- Logical masking, or the probability that the logical path's boolean equation may block an incoming SET, is obviously technology agnostic.
- Electrical masking, or low-pass filtering of narrow pulses, is quite hard to analyze: while gate delay is well understood and defines the likeliness of glitch transmission, a complete picture must also include how smaller transistors will induce different SET shapes. In a very recent study for instance [73], it was analyzed that the generated SET width and the logic delay vary similarly with reduced voltage, but for different physical reasons: a lower supply means that it is easier to induce an upset, and SETs last longer, since the restoring currents (or "drive strength") are diminished. In brief, when not only the voltage but transistor dimensions are also reduced, analyzing electrical masking is not trivial.
- Temporal masking, finally, can become an issue at advanced nodes: when operating at higher frequencies, the probability for an SET to reach sequential cells within their sampling window, becomes increasingly likely: omitting flip-flop setup, hold, and clock transition times, this probability is approximately $T_{\text{SET}}/T_{\text{CLK}}$ (with straightforward notation). Therefore, as early as 1997 [74] it was predicted that the logic SER could exceed the sequential SER for VLSI circuits in modern nodes. This is shown in Figure 1.24, where SETs become dominant at high frequencies, while the SEU rate remains flat.

Trends in fully-depleted transistor architectures

Technology downsizing is not necessarily continuous. The above discussions are mainly true of planar bulk, but technology leaps also occur; from the second half of the 2010s,

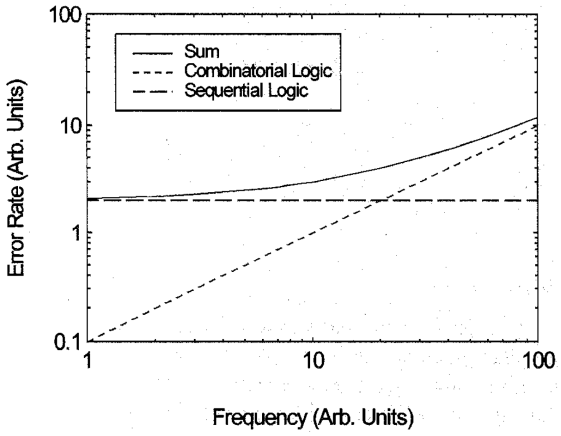

Figure 1.24: Single-Event Transient (SET) risks at high operating frequency. Taken from [74].

new CMOS transistor architectures have been proposed to overcome the limitations of planar bulk CMOS. The two main industry-adopted processes are FD-SOI and FinFET, whose electrical benefits have been largely praised: the reduced active silicon volumes in both technologies, thanks to the Buried OXide (BOX) in SOI or the vertically-grown "fin", leads to better electrostatic control of the (multi-)gate over the channel, thus better switching times and limited short-channel effects. Moreover, in both technology architectures the transistor is fully depleted, resulting in lower static consumption—unlike in bulk where free carriers in the substrate can induce leakage currents. Finally, variability is reduced, because substrate doping can be relaxed: both in FD-SOI and FinFET, the transistor's body is nearly intrinsic silicon (or silicon germanium).

Quite naturally, in terms of soft-error response, FD-SOI and FinFET have received a growing attention these past years. At the beginning of this PhD program, first measurements were disclosed, mainly focusing on terrestrial environments with alpha and

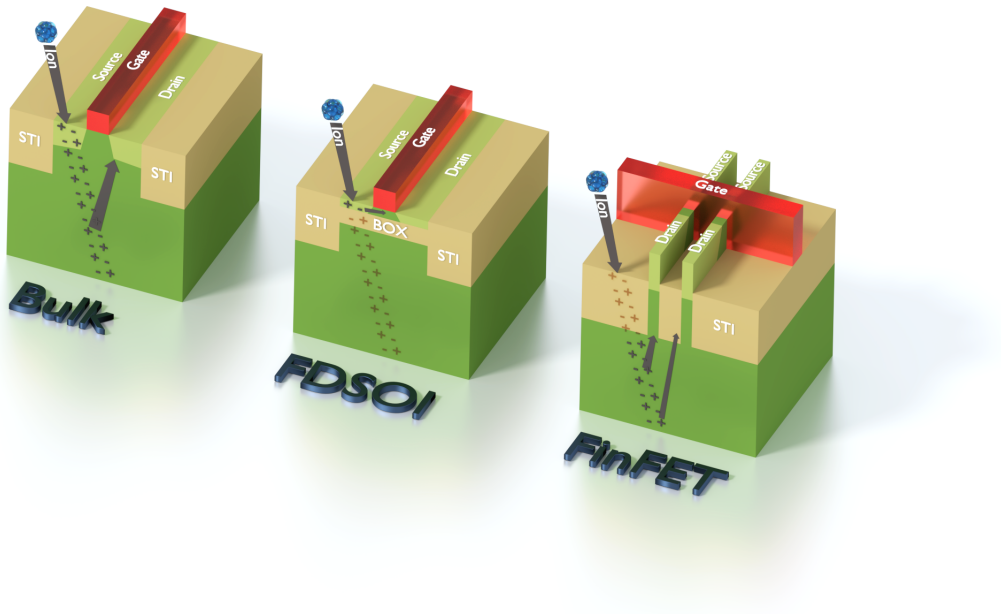

Figure 1.25: Charge collection mechanisms in planar bulk, UTBB FD-SOI, and bulk FinFET.

neutron irradiation [71, 75–77]. When compared to planar-bulk counterparts, sub-32 nm FinFET and FD-SOI display SER gains of 0.5–1 decade, and 1.5–2 decades, respectively. Please note that these are rough orders of magnitude based on the reported terrestrial particle test results. Although high-energy neutrons and protons yield similar silicon response, these are not enough to derive a cosmic-ray SER. The phenomenology for these trends is illustrated in Figure 1.25: in Ultra-Thin Body and Buried oxide (UTBB) FD-SOI, charges deposited in and below the insulating BOX cannot be collected—except perhaps from a thin BOX layer where electrons can escape at a very small, but finite mobility. Then in bulk FinFET, only a small fraction of the substrate charges may diffuse to the active regions. Now within the timeline of this work, more studies were performed targeting on-orbit environments (with heavy ions and high-energy protons), and highlighting the possible benefits for space technologies. This thesis largely subscribes to this orientation toward space applications, and additional, more recent data will be cited in this manuscript.

1.5.2 Space market evolution and new challenges

Beside microelectronic technological progress, recent evolutions in the space market ought to be mentioned. Up to this point, the historical milestones that were cited have been mainly in the pre-2000s discoveries of SEEs. However, a timely parameter that adds up to the context of this thesis is the recent bloom in the space industry landscape: historically, state space programs have targeted long missions (frequently above ten years) in Geosynchronous Earth Orbit (GEO) and also for explorations outside the near-Earth environment; starting in the mid-2000s, many non-governmental groups have been shifting the practices toward Low Earth Orbit (LEO), populated with numerous, smaller satellites of shorter life spans (e.g. five years). As can be seen in Figure 1.26, the number of secondary payloads sent to orbit each year has been on the rise, while the number of launches has kept somewhat even. Launch remains one of the main drivers, and obstacles, in space activities: it requires hundreds of people over months or years for proper operation, tens of millions or dollars or euros, and is by nature high-risk—due to sheer mechanical stress at high velocities, combined with fuel stability issues, safety margins in space flight are, when not negative, very thin and success rates of 95% are typically considered "good" [78]. Historically, this has led to very conservative designs and as far as the electronics is concerned, traditional space technologies have been lagging several nodes behind mainstream manufacturing processes. In short, many practices in the space industry can be traced back to the rarity, the expense, and the risk associated with rocket launches. This is still true today, and what Figure 1.26 evidences, is simply the fact that over the past few years, many more space missions have been designed to send smaller payloads to orbit, while not necessarily scheduling their own flight date or choosing their precise orbit.

The nature of these missions falls in two broad categories. First, universities see an increasing interest in projects with such short timescales and moderate costs. Indeed, many "SmallSats", or satellites below 180 kg, can be packed within a single launcher, thus reducing the launch cost per payload. "CubeSats", which are standardized SmallSats in units of $10 \times 10 \times 10$ cm^3, are being used by academics for a variety of experiments, for instance observation missions where LEO offers increased resolution compared to GEO. Radiation environments in LEO are not as aggressive as GEO, hence when combined with short mission times, and non-intolerable failure, such orbits are highly attractive. Typically, the on-board electronics use Commercial Off-The-Shelf (COTS) components,

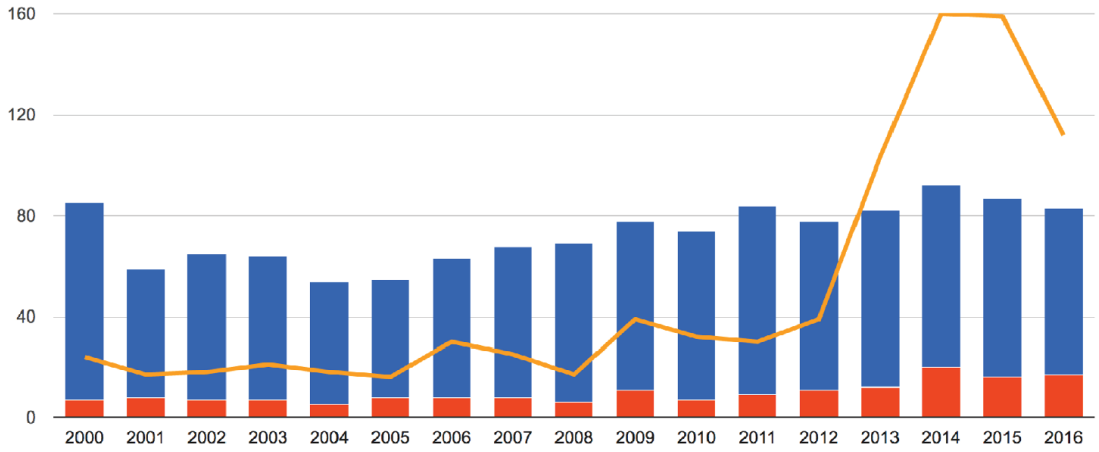

Figure 1.26: Yearly number of orbital launches without (blue) and with (red) secondary payloads, and number of secondaries launched (orange). Taken from [78].

and radiation effects are not necessarily a primary concern. Now, the other main class of organizations wishing to benefit from miniaturized commercial electronics, are private companies looking to commercialize the near-Earth environment: exomedicine or pharmaceutical studies in zero-g, space tourism, broadband Internet, are among the several applications being invested in [79]. Thousands of high-throughput satellites are expected to fly in outer space by 2025, with several satellite constellations aiming at global communication services [80]. These ambitious projects are all targeting LEOs, with low altitudes not only saving weight for more payloads (less rocket propellant), but also making it possible to achieve low latencies in the millisecond range—the characteristic drawback of LEO being the reduced coverage per satellite, hence the need for a fine constellation mesh. The targeted bandwidths of several gigabits per second also have direct implications for the embedded electronics: such throughputs are not achievable with conservative rad-hard chips manufactured in relaxed process nodes, and hence these projects are willing to harvest the benefits of technology downscaling in order to achieve the most cost-efficient performance–reliability compromises.

Of course, this is not to say that traditional space programs are a thing of the past: missions outside the near-Earth environment still need to be addressed, with the high levels of radiation tolerance that it entails, especially for explorations toward celestial objects whose radiation environment is unknown—or very sparsely quantified compared to Earth's. Nevertheless, with the applications mentioned above, commercial technologies are increasingly serving the space roadmap, and the technologies studied in this work fall within this scope.

1.5.3 STMicroelectronics' offer for space technologies

The first design and technology platform studied in this work is ST's C65SPACE platform, built around 65 nm vanilla bulk CMOS and further enabled for space applications with rad-hard features [81]. Beside electrical specifications [82] such as 1.1 ± 0.1 V operating voltages and -55–125°C functional temperatures, hardening is enforced at several levels: specific process modifications were made to ensure that the platform is free of SEL across the voltage–temperature range up to an LET of 60 MeV$\cdot$cm^2/mg. Then in terms of RHBD features, SEE waveforms are provided on demand for checks on sensitive nodes, and a comprehensive library of robust standard cells is proposed for SEU and SET hard-

ening. The IP offer includes silicon-proven rad-hard IOs, PLLs and High-Speed Serial Links (HSSLs), along with several memory compilers[7] for single- and dual-port SRAM and ROM with Built-In Self Test (BIST) and ECC features. Finally, aging models and TID qualifications allow to target twenty-year mission profiles.

Secondly, the principal object of study in this PhD is ST's 28 nm UTBB FD-SOI. SOI technologies have historically been confined to specific markets such as military applications requiring high levels of resilience to radiation. But nowadays, they are progressively becoming one of the industry standards for a variety of applications. As explained in Section 1.5.1, page 39, the scaling of the planar bulk transistor is reaching its limits, and ultra-thin SOI has been shown to enable faster, cooler transistors at low cost overheads (mainly due moderately-higher prices of SOI wafers compared to plain silicon). Using the Smart Cut method [83], it has become possible to manufacture SOI devices with a very thin silicon film sitting on top of a thin buried layer of silicon dioxide. Transistors in ST's 28 nm FD-SOI [84] feature a 7-nm channel controlled by a high-κ/metal gate (Figure 1.27). For SEE considerations, note that the epitaxy-raised source and drain implants, after contact formation via silicidation, lead to an active-silicon thickness between 10 and 20 nm in those regions. The BOX is 25-nm thick, which offers interesting back-biasing opportunities for either low-power or high-performance applications, respectively leveraging reverse and forward body bias. Since the channel is nearly intrinsic, most random dopant fluctuations are avoided, and the simpler implantation recipe also leads to fewer process steps which compensates some of the extra costs of the SOI wafer.

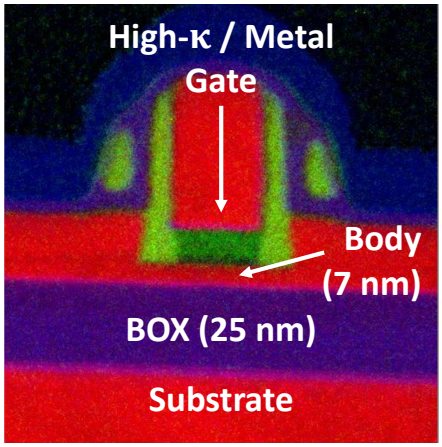

Figure 1.27: 28 nm FD-SOI transistor imaged via Transmission Electron Microscopy (TEM).

As far as radiations as concerned, anticipating a little on what comes next, it was already mentioned in Section 1.5.1 that FD-SOI offers excellent soft-error performance compared to planar bulk: the reduced silicon volumes lead to very small error rates, with MCU issues nearly eliminated since individual devices are isolated from each other thanks to the BOX. While this is "good news", it is also incredibly challenging for the experimentalist since very few events are captured during test campaigns. This intrinsic

[7] A memory compiler is a generator that allows to create the Electronic Design Automation (EDA) database—incl. physical layout, schematic, logic-gate and Register-Transfer Level (RTL) netlists, Place & Route (P&R) view—for a memory cut of chosen configuration—e.g. words, bits, multiplexing factor, number of banks—and made up of a certain bit-cell. In other words, it is a piece of software that creates the memory array and read–write periphery for a cut of arbitrary dimensions.

soft-error tolerance for ground-level applications, however, is not sufficient for space missions as is. For that reason, ST has developed a 28 nm FD-SOI space platform leveraging on commercial FD-SOI [79]. Note also that the BOX natively provides SEL immunity, something that required paramount care in bulk technology.

1.5.4 Research objectives

In light of all the space environment notions, radiation physics, modern CMOS trends, and recent industrial evolutions that have been discussed along this chapter, the research objectives of this work can now be stated. As was already underscored, this thesis is firmly oriented toward SEE modeling, with a main focus on 28 nm FD-SOI technology:

- A primary topic is to comprehend the charge transport mechanisms that are specific to advanced UTBB SOI. This implies the use of TCAD simulations as an entry point to delineate the characteristics of the technology under radiation, and is justified by the industrial interest in FD-SOI for space applications. Although general literature exists on SEEs in FD-SOI, particular charge collection mechanisms pertaining to 28 nm FD-SOI need to be carefully studied.

- Another scientific goal lies in the development of compact SEE models in FD-SOI, to be used for SER prediction at circuit level. Physical insight gained from TCAD investigations must be transposed into numerically-efficient methods to compute radiation-induced currents at minimal accuracy loss. This part of the work is where lies the most custom physics in this research.

- A more technical, but directly related aspect is the upgrade of ST's Monte Carlo simulation tool TIARA, in order to enable SER prediction in FD-SOI. With charge-collection models specifically tailored to FD-SOI, plug-and-play integration in a simulator originally designed for planar bulk architectures is not possible. Several software developments have to be undertaken in order to maintain TIARA at state-of-the-art level for fast, accurate SER predictions in fully-depleted technology, with the maximum level of automation and integration within the standard CAD flow.

- Finally, the research and development tasks above must support the space industry revolution, i.e. by using high-end technologies such as 28 nm FD-SOI with short time-to-market. In that context, modeling tools must serve to optimize analysis times with regard to silicon iterations and experimental campaigns. This means projecting the cosmic-ray SER response of modern CMOS technologies—namely in 65 nm bulk Si and 28 nm FD-SOI—in order to assess circuit sensitivity and perform design optimizations prior to tape-out, and more generally to draw generic rad-hard guidelines to consolidate these space platforms. This is all the more so true in FD-SOI, where the ultra-low intrinsic sensitivity leads to very sparse measurement data on which to base the design choices and hardening strategies. Exposure times must be increased to collect satisfying statistics, hence less environments can be tested—e.g. not all energies and angles can be scanned; at constant silicon area for the irradiated test vehicles, lesser cell types can be selected, and only a subset of the libraries can be directly qualified. In this context, fast and accurate simulations are expected to fill the gaps (interpolation) within the experimental datasets, and enable SER projections (extrapolation) beyond the observability of standard test procedures, to accommodate the orbit environments of timely applications in LEO and GEO.

Chapter 2

TIARA: a Monte Carlo soft-error rate simulation platform

In this second chapter we introduce TIARA, STMicroelectronics' and IM2NP's SER simulator at analog-circuit ("SPICE") level. After giving a succinct history of SER simulation at ST, we highlight the salient features of TIARA and its software structure and code characteristics. We then move onto a more detailed description of its subcomponents and assets. Note that the primary goal here is to put TIARA's software-engineering aspects in the spotlight but inevitably, several of the code design choices are intertwined with the underlying physics, which we thus also discuss ever so briefly.

2.1 Overview of TIARA

2.1.1 A brief history of soft-error rate simulation at ST

Monte Carlo SER simulation at logic-cell level has been a continuous research activity at ST since the early 2000s, in collaboration with several academic partners. Contributions have been somewhat heterogeneous over the years, owing partly to the challenges of coordinating R&D efforts across various PhD programs. Nevertheless, the emerging logic is that major upgrades to the simulator have always been impulsed by disruptive changes brought to its core engine, i.e. the charge transport and upset models:

- An initial simulator was developed in [85], which modeled parallelepiped sensitive volumes and computed upset rates based on a critical-charge criterion—in other words, counting SEUs for events whose associated Q_{dep} in the sensitive volume exceeded Q_{crit}. Even though the considered Q_{crit} was static, this allowed to study the intrinsic distinction between bulk Si and SOI technologies, based on different sensitive silicon depths.

- Then based on original work in [7], in the mid-2000s [86], a model was implemented where for each impact, the transient radiation-induced current was computed via an ambipolar diffusion equation. Then the transient's peak coordinates were checked against an $I_{\max} - T_{\max}$ critical curve to set the SEUs apart from the harmless events. This allowed to capture the time dynamics of upset mechanisms, however this critical curve had to be precomputed in SPICE and the diffusion model applied primarily, if not exclusively, to bulk Si technology.

- In [5], systematic SPICE calls with radiation current injections were introduced to avoid circuit precomputations, and the current pulse model was also improved to better account for drift currents (as opposed to diffusion). These late-2000s developments to the simulator were made alongside with its usage across a variety of standard and rad-hard cells down to the 32 nm bulk Si node. This truly gave rise to an industrial tool set, and this is when the name TIARA was coined. In 2014 an anthology article was published [67], covering the different physical blocks in the simulation chain and the various applications that the simulator had allowed to serve.

- Now in the framework of this mid-2010s thesis, as we will abundantly discuss in the next chapters, the central scientific additions made to TIARA are the developments of collection models for FD-SOI technology architectures, down to the 22 nm node. In terms of engineering work, massive overhauls were also undertaken with the tremendous help of software specialists to improve the tool's usability and scalability. Although they may be regarded as less challenging from an R&D perspective, they are the 'D' in "R&D" and should certainly not be overlooked.

2.1.2 TIARA "datasheet": main features and key figures

A virtual irradiation experiment in TIARA is always defined as the combination of a target circuit, under a certain environment of incident radiation; therefore the simulator's principal traits can be summarized by the scope of radiations and circuits it can handle.

In terms of particles first, TIARA can natively transport ions across the whole periodic table (atomic nuclei with their stable isotopes) thanks to pre-generated SRIM LET curves in all necessary target materials. For nuclear interaction of protons and neutrons with silicon, TIARA relies on external databases built from Geant4 to generate the resulting ionizing products. This gives the tool the ability to account for virtually any environment ranging from sea level (mainly alphas and neutrons) to space orbits like LEOs (nearly just protons) or GEOs (primarily protons but also heavy ions). Subatomic particles such as electrons and muons—of which high fluxes are encountered in the Van Allen belts and on Earth, respectively—are not treated however: their contribution to the overall error rate in digital ICs is, if not negligible, at least not dominant for the technology nodes in this work [87–90].

Now in terms of electronics, we shall say that TIARA is primarily intended for the simulation of CMOS digital ICs, even though the modeling concepts used may apply to a variety of technology processes and circuit architectures. Thus typical simulation setups are meant to study SEUs in memory or sequential logic cells, or SETs in combinational logic cells. TIARA supports planar bulk Si down to the 32 nm node, based on charge transport in the substrate and collection from the MOSFETs' implant junctions—a logic which could be extended to handle bipolar devices. TIARA also handles planar FD-SOI technology down to the 22 nm node, based on charge transport and collection in confined silicon volumes—a logic which could be extended to handle FinFET devices which are also *fully depleted*.

Having said this, a few figures can help illustrate the perimeter of SER simulations in TIARA: in terms of simulation scale, it can typically handle a few dozen devices in the physical domain, i.e. the 3D structure within which radiation calculations are carried out. In the electrical domain, TIARA being in a sense no more than a very elaborate wrapper to an external SPICE program, its scope is inherently set by the limitations

of analog circuit solving: a few hundred (non-trivial) devices are a reasonable setup, and for example running several thousands of transistor instances means pushing the envelope. For typical simulation setups, this all leads to computation times of 1–10 seconds per *fully-processed* impact. Indeed, many events can actually be *pruned*, which means that TIARA avoids going through the full processing chain for instance when the total deposited charge is negligible, or when the struck devices have little voltage difference on their terminals. This, and the fact that TIARA can operate on a computer cluster (typically between 10 and 100 jobs in parallel), yields simulation times ranging from about half an hour to half a day: a simple heavy-ion scan (a few thousand impacts per species) on a bit-cell (a single electrical state to test for a symmetrical SRAM) can run in ten minutes; a full-spectrum proton irradiation (billions of protons to trigger hundreds of thousands of nuclear events) on a flip-flop (8 possible electrical states for a master–slave DFF) will typically run overnight and provide fresh results in the morning. Now as far as accuracy is concerned, in general for SER estimation, making predictions within 2× w.r.t. experimental measurements is very reasonable, considering the complex interplay of physical mechanisms involved. In TIARA this is what can be expected from, say, extrapolative simulations on new technology nodes or projections far outside the calibration domain in a given technology. When running with adjusted models on a well characterized process, the simulator typically achieves ±30% accuracy when compared to beam test data.

2.1.3 Software architecture

TIARA's task is to evaluate the soft-error sensitivity of a given circuit under a particular radiation environment. To that end, it makes use of the Monte Carlo method, which means performing random sampling on the input variables—in our case, the incident

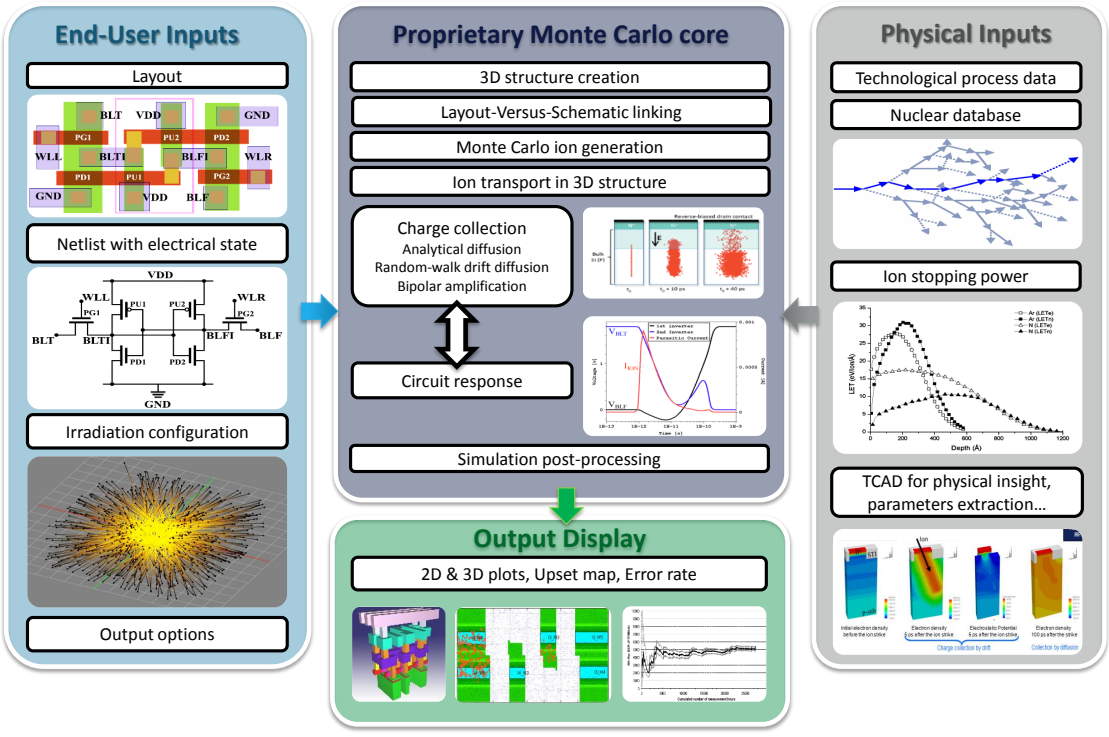

Figure 2.1: Architecture diagram showing TIARA's main components and inputs–outputs.

particle's parameters—to compute a numerical integral—in our case, the circuit's radiation cross section or upset rate. More specifically, like most Monte Carlo algorithms, TIARA operates in three steps: sample, evaluate, and aggregate. Or said differently, to find the integral of a given $f(x)$, first pick random x values, then evaluate function f at these points, and then add up all the $f(x)$ values. This differs from deterministic integration schemes only in that the x are random variates instead of deterministic numbers, such as for instance evenly-spaced points on a grid. The mathematical implications, advantages, and shortcomings of this are discussed at greater length in Appendix A. Having said this, we can describe TIARA's main procedure in terms of a Monte Carlo loop that repeatedly draws random particle events (the "sampling" phase), and for each of them, computes the silicon's response to the radiation-induced currents (the "evaluation" phase) and gathers statistics that finally result in macroscopic quantities such as the circuit's error rate (the "aggregation" phase).

Having set the scene, we can now consider Figure 2.1 which outlines TIARA's software architecture and high-level modules.

- Given a layout file (or IC mask design) of the circuit, a so-called `Builder3D` module (Section 2.2.1) constructs a three-dimensional structure of the circuit in the physical domain. The module also takes as input a geometrical description of the technology process stack: roughly speaking, the layout provides us with the *top-view* data ('X&Y' coordinates), and the process stack is the *side-view* information needed to extrude the patterns into layers of a certain thickness along the 'Z' axis.

- The user must also provide a transistor-level netlist of the circuit to simulate, which will be used to compute the electrical response to particle strikes. Therefore, a physical-to-electrical mapping is needed for TIARA to identify victim *circuit nodes* to inject the radiation-induced currents at, when provided only with the *geometrical location* of the event. This is performed by the `LVSmatcher` module (Section 2.2.2) , which runs a custom Layout Versus Schematic (LVS) comparison to match the layout and netlist connectivity graphs. Together with the `Builder3D`, these modules form the preprocessing stages that run before entering the Monte Carlo loop.

- Now inside the Monte Carlo loop, the "sampling" phase mentioned before is the procedure of picking the positions, directions, energies and masses (i.e. nucleus composition) of particles striking the considered circuit, based on a description of the environment. In other words, the `Irradiator` module (Section 2.3) draws random variates from probability laws that describe the statistical distribution of the incident particles' properties, which are both geometrical and spectral (e.g. with energy-dependent fluxes).

- Then the "evaluation" phase, which ultimately must determine the response that each event triggers on the circuit, roughly follows the physical processes that were described in Table 1.1 on page 11: first, a module named the `Raytracer` (Section 2.4.1) is called to compute the initial charge deposition induced by the particles that the `Irradiator` generates. It does so by intersecting particle rays with the geometry from the `Builder3D`, and computing the associated energy losses along the resulting segments based on tabulated stopping power curves.

- Following charge-deposition calculations in the `Raytracer`, the `Collector` module (Section 2.4.2) is in charge of running the semiconductor physics responsible for

the transport of charges in active silicon, and their collection at electrical contacts. These processes will be discussed in Chapter 3, together with the associated modeling challenges—which, in brief, amount to preserving the right physical trends whilst meeting severe speed requirements, namely being able to run an impact in no more than a few seconds rather than minutes or hours as with TCAD simulators. Note that while all other modules in TIARA are entirely technology agnostic, different versions of the `Collector` module exist, suiting various needs in terms of speed–accuracy trade-offs or technology architectures (e.g. planar bulk Si and FD-SOI call for radically different modeling choices and physical assumptions). In a sense, this is where in TIARA lies the most *custom* physics, and where most of the scientific efforts were concentrated throughout this thesis.

- To evaluate the circuit's response to currents triggered by the particles, the `Collector` is coupled to a `CircuitSolver` module (Section 2.4.2) which inserts time-varying generators in the simulation netlist, and encapsulates system calls to an external SPICE solver left in charge of computing the circuit transient evolution under those stimuli. The `Collector`–`CircuitSolver` link can either be sequential or fully dynamic, depending on whether the radiation currents are computed considering a *frozen* electrical state *right before* injection time, or computed *self-consistently* with the circuit's evolving state (see remarks on overlapping device–circuit response times on page 11; the importance of circuit coupling for collection models will be further discussed in Chapters 3 and 4.).

- Finally, the `Analyzer` module (Section 2.5.1) performs post-processing steps on the electrical simulation outputs to assess whether single-event effects occurred: current and voltage waveforms obtained for the simulated event are compared to a *reference* simulation, meaning a golden run without radiation injection. For instance, this allows to flag an SEU if there is a significant voltage difference at final time on a bistable node, or measure the Full Width at Half Maximum (FWHM) of SETs in logic paths by looking at how long the voltages swing away from the golden run. These *micro-analyses* are then also used by the `Analyzer` to draw statistics upon and derive *macroscopic* quantities; this corresponds to the "aggregation" phase, where we compute for example the cross section or the error rate knowing the overall fluence or flux generated by the `Irradiator`.

- This whole chain of processes is orchestrated by a higher-level module named the `Driver` (Section 2.5.2), which controls the sequence in which all modules listed above are run: first the preparation stages involving the `Builder3D` and `LVSmatcher`, then the Monte Carlo algorithm itself which repeatedly loops over the `Irradiator`–`Raytracer`–`Collector`–`CircuitSolver`–`Analyzer` chain. In other terms, the `Driver` is our simulation *harness*, which ensures seamless execution of the subroutines with all the exception handling that it entails, whether TIARA is running in standalone mode—single-machine execution e.g. on the user's workstation—or in parallel mode, i.e. when calculations are distributed on multiple nodes.

2.1.4 Technical characteristics

TIARA is written in C++, in about 53 000 lines by today's count. The submodules introduced before are implemented as classes (in the sense of object-oriented programming), and the TIARA program itself is an executable binary whose `main()` function

instantiates the `Driver` and lets it run. Surrounding this core code are Python and shell scripts to automate certain tasks before, during, and after simulations. To list a few, these scripts allow to generate some configuration files and input setups, to monitor jobs when running in distributed mode, and to post-process raw outputs for high-level summaries or visualizations. One of our main external dependencies is the Boost library [91] for several generic services it offers (e.g. to abstract some system calls, or for option parsing on the command line) and more importantly, for 2D geometrical operations required when manipulating layout data in the `Builder3D` and `LVSmatcher` modules. We also make use of Eigen [15] for matrix operations and linear algebra required for semiconductor solver calculations in the `Collector`. Finally, for 3D geometry operations throughout all modules we rely on J5geolib, a library developed at IM2NP.

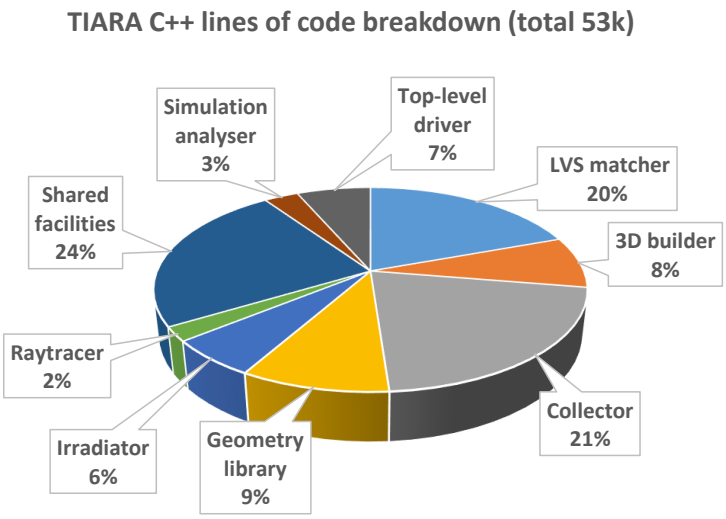

Figure 2.2: TIARA's C++ source code broken down between its constituents.

Looking at Figure 2.2, we see that beside "shared facilities" (internal data structures describing radiations and circuits, and input–output APIs), the `Collector` code is the largest slice in our lines-of-code pie chart. This reflects the very custom nature of the compact models implemented for charge transport. The preparation stages (`Builder3D`, `LVSmatcher`) also represent a fair amount of the code, due to the numerous particular cases encountered—input circuit files can never be completely harmonized, being generated from sometimes heterogeneous CAD flows. These modules for circuit structure creation may not perform any physics but they are what truly sets the industrial dimension in the tool: at the time of this thesis, TIARA has been used for dozens of SEE studies and simulation structures have been successfully generated for hundreds of logic cells. Finally, the Monte Carlo processing blocks outside the `Collector` (`Raytracer` and `Irradiator`) are relatively thin, but this may be somewhat biased by the fact that they heavily rely on J5geolib ("geometry library" in Figure 2.2), while other modules rarely call J5 geometrical routines.

TIARA's input–output system is largely text-based at the moment. Simulations are launched in a terminal via a command line to call the `Driver`'s binary with options defining top-level inputs such as the circuit to simulate, the radiation environment, and the required analyses. We use XML files to configure TIARA's submodules—which proves to be very flexible, albeit a bit "verbose". We support Graphic Database System (GDS) format (binary files) for layouts. As for SPICE netlists (text files), we are able

to parse from Circuit Design Language (CDL) and Detailed Standard Parasitic Format (DSPF) formats respectively for pre- and post-layout netlists, i.e. without or with RC parasitics extracted from the physical design. Other than this, physical inputs stemming from external programs (SRIM and Geant4 databases) have specific text formats: these are the ones that required developing custom parsers, while for XML, GDS, and CDL–DSPF inputs we rely on free, open source libraries under GPL. Now as far as outputs are concerned, again we mostly generate text files: high-level simulation outputs are dumped to a text summary file, and per-event result databases are recorded as Comma-Separated Values (CSV) files where each event registers several (customizable) fields: this results in a table with as many lines as events, and for instance two columns for the impact '$X\&Y$' coordinates, and another column for the SEU boolean flags—a raw file which we can then be converted to an "error map" image via a dedicated Python script. 3D outputs are also generated in text following the OBJ (Wavefront) format specification [92], and opened in external tools such as Blender [93] for visualization.

Lastly, in terms of platforms and infrastructures, let us mention that TIARA commonly runs on Red Hat Linux machines, and occasionally it has been compiled for other Unix-based systems like MacOS. Beside the operating system, more stringent requirements are imposed by the fact that currently we only support Load Sharing Facility (LSF) from IBM [9] for distributed execution, and Eldo from Mentor Graphics [8] for the SPICE simulation block. Note however that HSPICE [35] netlists (Synopsys) can also be run thanks to Eldo compatibility modes.

2.2 Circuit structure creation

For a TIARA simulation to run, first of all we create a structure that represents the circuit both in the physical and electrical domains.

2.2.1 Physical-domain modeling: the `Builder3D` module

The `Builder3D` module first reads the layout file, which provides the $X\&Y$ coordinates of the 2D patterns drawn on each of the CAD layers. TIARA then combines the CAD layers using boolean operations to create physical geometries. For instance, the transistor gates are defined where polysilicon bars are *over* active silicon areas (*intersection* operator), and the STI regions are the *complement* of active silicon (*not* operator), as illustrated in Figure 2.3a. This amounts to performing CAD-to-mask calculations, only quite simplified: we do not account for rounded geometries arising from finite lithography wavelengths (and the associated compensations at CAD-to-mask time) or chemical etching "smoothness", for example. TIARA's 3D structures are collections of axis-aligned boxes, which seems a reasonable choice considering the Manhattan-only geometry for the layouts in modern technologies. Note that layout files encode arbitrary polygons (closed chains of vertices), and thus TIARA must first perform a meshing operation to convert the paths into non-overlapping rectangles. For all these geometrical calculations (boolean operations and meshing), we make great use of routines available in the Boost library. At this point, we also account for the optical shrink of the technology if any, by multiplying the $X\&Y$ coordinates by a certain factor, e.g. 90% for ST 40 nm designs drawn as in 45 nm.

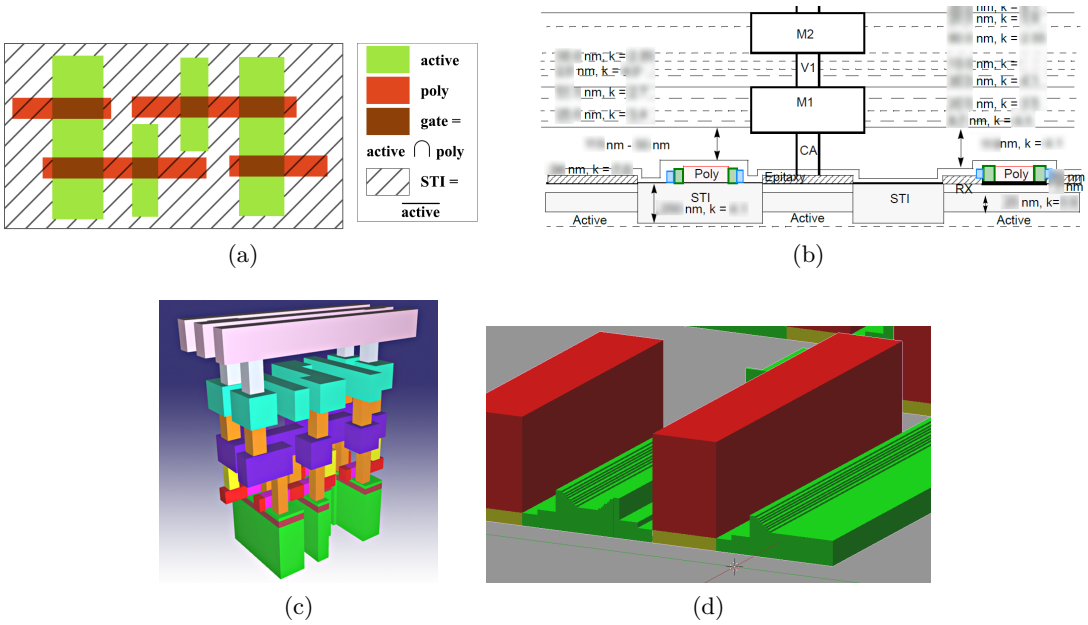

Figure 2.3: `Builder3D` inputs and treatments to create a simulation structure: (a) Boolean operations on the layout layers to derive the physical masks – (b) Snapshot from an ST Design Rule Manual (DRM) document feeding the TIARA "process stack" file with thickness and material data – (c) 3D structure obtained after extrusion for an FD-SOI SRAM cell (InterMetal Dielectrics (IMDs), STI and substrate beneath not shown for clarity) – (d) Transistor close-up with detailed SOI thickness modeling.

Process stack layer	**Typical materials**	**Approximate thickness**
Substrate (wafer bottom)	p-type Si	1 mm
Wells	n- and p-type Si	1 µm
Implants	n^+ and p^+ Si	100 nm
STI	SiO_2	200–300 nm
BOX in FD-SOI	SiO_2	20–30 nm
Active film in FD-SOI	Si(Ge)	5–20 nm
Gate dielectric	high-κ e.g. HfO_2 (≈ sub-90 nm); SiO_2	a few nm
Contact pillars	W	50–100 nm
Vias, metal lines (bottom)	Cu (≈ sub-90 nm); Al	100 nm
Vias, metal lines (top)	Cu (≈ sub-90 nm); Al	1 µm
IMDs	low-κ (≈ sub-90 nm); SiO_2	Same as interconnects
Passivation	SiN, SiO_2	1–5 µm

Table 2.1: Typical data contained in TIARA technology process stack files. See Figure 1.11 on page 20 for an illustration of the respective layers.

After the 2D geometry is generated, we extrude it on the Z axis (see Figs. 2.3b–2.3c) with thicknesses and materials information described in a specific TIARA "process stack" file, with typical values given in Table 2.1. This file is populated with data collected from the technology DRM, TCAD structures, or Transmission Electron Microscopy (TEM)

cuts for instance. At this point, we also group some of our 3D boxes in containers, mainly to accelerate `Collector` calculations later on: boxes belonging to a same transistor implant are clustered as such, and connected boxes of n-well, p-well, p-substrate, are grouped by island. This is done via a generic graph partitioning algorithm that iteratively grows sets of abutted boxes with common material properties. As shown in Figure 2.3d, for FD-SOI technologies with very thin sensitive volumes, we also take care in modeling the detailed morphology of the silicon film, to account for the varying thickness near the gate spacers and the silicided contacts reducing the height of the epitaxy-raised source-drain implants. This special operation cannot be performed with geometrical boolean operations, and is handled by adding geometry on both sides of the "channel" box. Another exception to the "boolean-operations + extrusion" paradigm is the "poly-bias" option, whereby transistor gates are incremented from their CAD-drawn length—this allows circuit designers to derive several standard cells from the same master design database, in order to serve different speed or power purposes. Similarly to the detailed SOI thickness modeling, the `Builder3D` handles poly-bias by retargeting the gates and channels of the MOSFETs along their length axis.

At this point, the TIARA circuit structure thus represents 3D geometry and materials from the bottom of the silicon wafer to its top after all the front-end fabrication steps, i.e. up to the passivation layers. We do not model the chip package (back-end fabrication) for practical reasons:

- While on orbit, most of the shielding is provided by the spacecraft itself and not the chip packaging—with the particular exception of antenna panels or solar arrays; environment modeling tools such as CRÈME or SPENVIS are able to provide apparent particle fluxes behind shielding and therefore, TIARA neither needs to transport particles in the package materials nor in the satellite structure—this is more of a concern for the satellite designer than for the silicon foundry.

- TIARA is also frequently used to simulate beam test situations—be it for model calibration purposes, i.e. comparison between experimental and simulated data on an already-tested circuit, or for direct comparison between, say circuit A with available experimental data and circuit B without experimental data but simulated with the same beam test conditions. As far as accelerated testing goes, the range of high-energy protons is great enough that chip packages need not be opened, thus the effective spectrum in the region of interest is only weakly altered. Hence TIARA should only incur small errors by not considering the wafer surroundings during high-energy proton irradiation—on the contrary, modeling the BEOL metal lines and dielectrics is highly relevant, since local shadowing effects may arise for short-ranged secondaries generated in the vicinity of sensitive silicon. The same would go for (high-energy) atmospheric neutrons—whose range is even higher than that of protons—if we were to consider terrestrial applications.

- Now for heavy-ion tests, where space-relevant LETs are achieved with energies quite lower than actual GCRs, the chips must be delidded so the ions can reach the Front End Of Line (FEOL). This once again means that TIARA can safely ignore the packaging layers when running heavy-ion beam simulations.

- To complete the picture, let us mention that α-rays from radioactive decay have very short range in solid materials (a few microns). Thus, whether they are emitted from the package materials in real life, or from an alpha source during accelerated

testing (with the chip package removed), the effective emitting region can always be modeled as a thin layer on top of the wafer.

To sum up, careful case-by-case analysis justifies the choice of not modeling the circuit structure beyond the wafer itself in TIARA, whether for real or accelerated environments, for space or terrestrial applications.

2.2.2 Electrical-domain modeling: the `LVSmatcher` module

After the geometrical structure of the circuit is built, the `LVSmatcher` module is used to map physical locations to electrical nodes in the netlist. This is done in three steps: we first parse the netlist to construct the connectivity graph of the schematic, by growing or updating its adjacency matrix for each device we read: transistor terminals can either declare a new node in the graph, or add an edge between known nodes. Likewise, the connectivity of the layout must be extracted, which once again involves several geometry routines from Boost: electrical nodes are propagated along polygons from the same layer touching each other, or through layers when a vertical connection (or "via") overlaps the geometries.

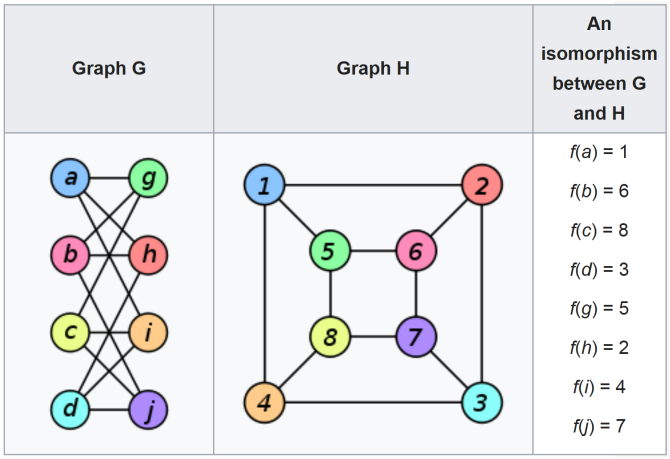

Figure 2.4: Graph isomorphism illustration showing two identical graphs in terms of connectivity, despite their different looking drawings. In the `LVSmatcher` the graphs being matched are the layout- and the schematic-extracted connectivities. Taken from [94].

Finally, the actual matching operation amounts to a known problem called the *graph isomorphism* problem (Figure 2.4): in order to decide whether two graphs are identical from a connectivity standpoint, oftentimes algorithms try to construct the correspondence between them. For this, we make use of a generic back-tracking algorithm available in Boost. To reduce the number of iterations, we guide the algorithm with as much *a priori* information as possible: for instance, transistors with different length or width (known from the netlist), or number of neighbors (their "degree" in the two graphs) can never be matched. This is forced by declaring them with different "vertex invariants" (a vertex invariant is a quantity that remains unchanged in the graph correspondence). This is absolutely crucial for the back-tracking algorithm to work, since it has exponential complexity. LVS-dedicated algorithms can perform faster, by iteratively refreshing the vertex invariants based on the devices' neighbors, then neighbors' neighbors, and so forth. Nevertheless, TIARA's in-house LVS has allowed to construct structures for hundreds of logic cells from varied standard cell libraries (inverters, buffers, standard,

high-speed, or rad-hard flip-flops), memory offers (single- or dual-port cells), and portions of custom designs in several technologies, such as level shifters or custom latches. Limitations may arise when trying to simulate larger structures involving more than just a handful of cells. To that end, the `LVSmatcher` module can be bypassed if the user directly feeds the LVS correspondence to TIARA as a text file containing the layout locations of the schematic MOSFETs. A future improvement will be to directly read this information from an LVS report generated from industrial tools.

2.2.3 Resulting data structure: the Circuit class

The end-result of the `Builder3D` and `LVSmatcher` is a simulation-ready structure, organized as shown in the Unified Modeling Language (UML) class diagram[1] of Figure 2.5.

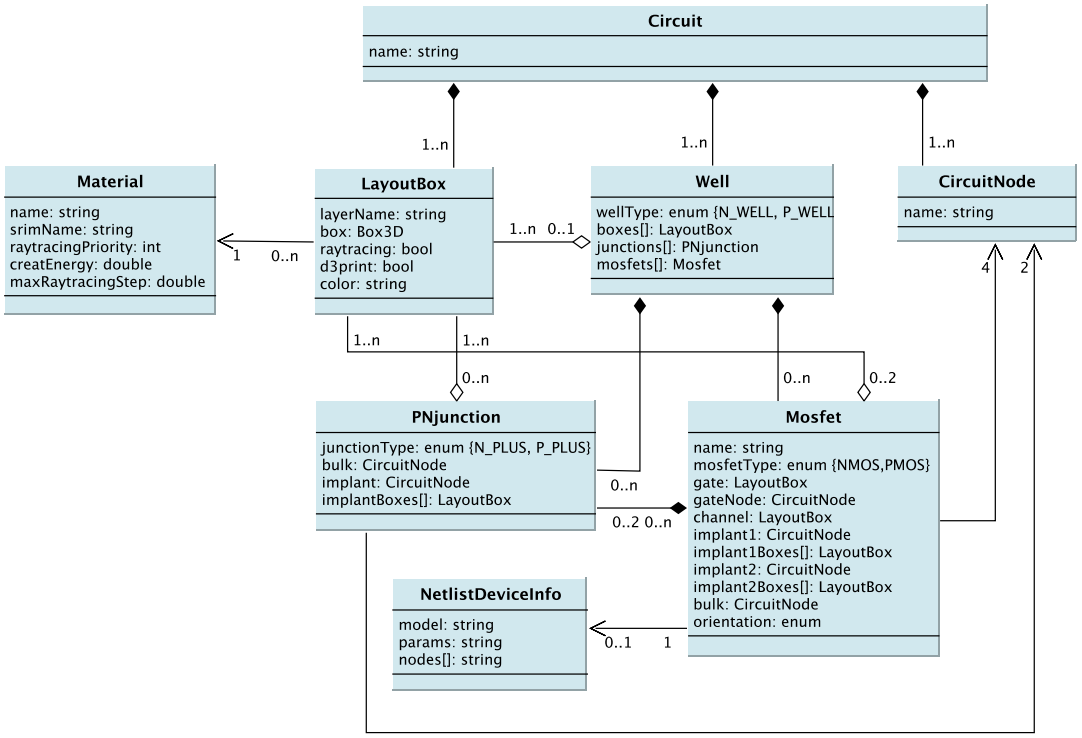

Figure 2.5: UML class diagram for the circuit data structure. Regular arrows indicate *association*, i.e. class attributes (member variables) with multiplicity; diamond arrows symbolize a set/element relationship, either with weak or tight coupling: *aggregation* (white diamonds) is typically implemented with pointers, while *composition* (black diamonds) is typically implemented with values so as to subordinate the life cycle of the elements to that of the set.

An instance of the Circuit class contains the 3D model of the circuit including its materials, for the `Raytracer` to transport ions inside; for the `Collector` module to transport electron–hole pairs in bulk technology, the Circuit class bears the information associated to the wells that confine carrier diffusion; the wells (and the p-substrate) enclose p-n junctions located at the transistor implants (n^+ over p-well or p-substrate and p^+ over n-well), whose mostly-vertical electric field can separate the carriers and lead to radiation currents; in FD-SOI technology, the basic blocks used by the `Collector`

[1]Note that the class diagram of Figure 2.5 is a static representation of the data structures manipulated by the modules. It is not showing the dynamic interactions of the modules with each other during the execution flow, which would be the goal of a UML activity diagram instead.

are directly the MOSFETs with a mostly-lateral field, since we ignore transport below the BOX. Finally, the `CircuitSolver` module makes use of these sensitive devices and of circuit node objects to perform the necessary current injections. Overall, today's Circuit class in TIARA is similar in essence to what was present in [5] and associated articles [66, 95]; the huge step-up rather lies in the fact that our Circuit objects are now created fully automatically thanks to the `Builder3D` and `LVSmatcher` flow, when previously only a handful of architectures were supported—3D volumes for SRAMs and DFFs were passed via specific text files instead of reading generic GDS files, and the electrical correspondence was hard-coded altogether rather than performing an LVS to find the mapping. This industry-oriented effort of software engineering has shown to greatly improve the throughput of our simulation set-ups, thereby vastly extending our spectrum of possible analyses.

2.3 Monte Carlo particle generation: the `Irradiator` module

We shall now move our focus onto the submodules running during TIARA's Monte Carlo loop to process radiation events. The first step is to randomly generate the ions incident on the circuit, which is the task of the `Irradiator`.

2.3.1 Path-finding approaches and flow overview

Mathematically, this means nothing but drawing random variates (the particles' energies or incident angles for instance) from certain probability distributions that are determined by the irradiation scenario we need to simulate. These specific distributions are achieved by picking random numbers from a *uniform* law, and then applying transformations on them in order to yield the wanted probability density. In Appendix A we further elaborate on how to and why apply such transformations to obtain the desired distributions. In any case, the cornerstone routine behind the `Irradiator` is to generate uniformly-distributed random variates between 0 and 1 and for that, we use a popular Pseudo-Random Number Generator (PRNG) called the Mersenne Twister [96,97], which we also quickly discuss in Appendix A.

In [5], simulation studies were carried out with alpha particles, neutrons, and *unidirectional* heavy ions mimicking beam tests. For all of these, the radiation event generation can be done by first picking a starting point inside the emitting source, and then assigning a direction to the ion(s)—neutron simulations making use of nuclear reaction databases follow that same pattern, with the emitter being the silicon itself, in a way made radioactive by the incoming neutron flux. Now, this thesis being geared towards space applications, the irradiation scenarios we also have to cover are on-orbit situations where the heavy-ion fluxes can be *omnidirectional*, and proton scenarios for which the flow is identical to neutrons, as long as Proton Direct Ionization (PDI) can be overlooked (see Section 4.3 on page 121 for a more thorough discussion).

For on-orbit omnidirectional heavy ion scenarios, a difficulty arises when trying to sample the events using the procedure described above: if we construct rays by choosing their *starting* point inside a finite emission *volume*, we introduce a bias by not considering emission points further away: even though we may very well reach all possible locations on our targets, the statistics of the impacts will be wrong unless the emission volume is made very large. This can be overcome by sampling the ray properties the other

way around, from target to source instead of going from the source to the target. In such a *backward* path-finding approach, first the *ending* point of the ray is sampled on the *surface* of the target, and a direction is chosen. Then only, the starting point is deduced by "rewinding" the particle outside of the 3D structure before sending it to the `Raytracer` for energy loss calculations.

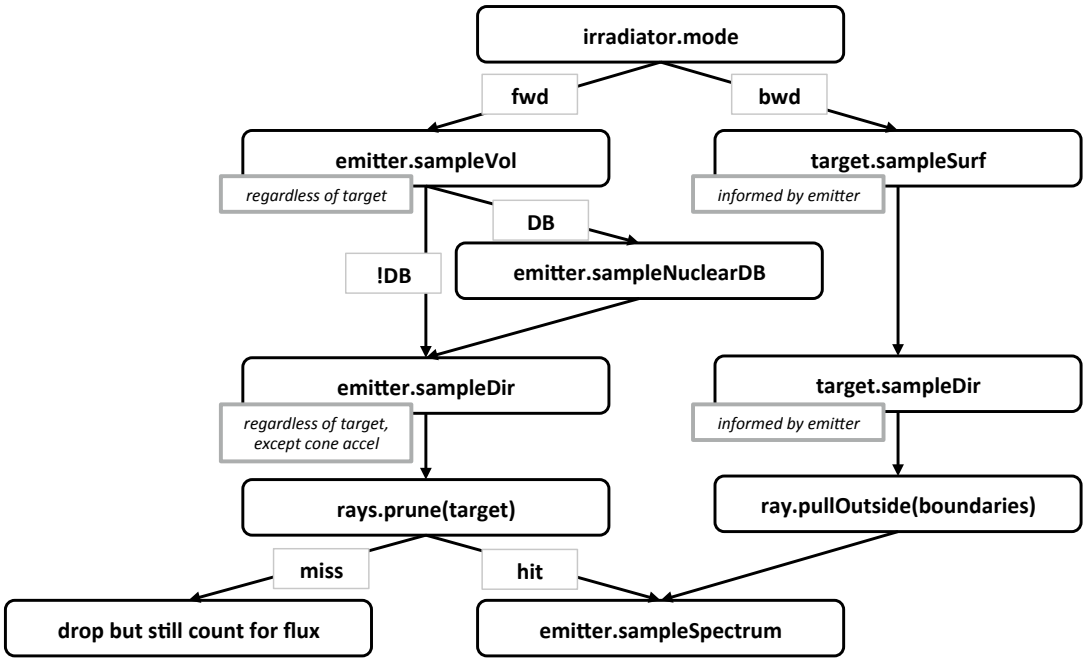

Figure 2.6: Flow diagram of the `Irradiator` module.

For that reason, the `Irradiator` was redesigned to offer both forward and backward path-finding modes in a unified and flexible architecture, as depicted in Figure 2.6. Irradiation scenarios are modeled with three main data structures, namely an emitter, a target, and a "boundaries" object. They encode for specific properties of the `Irradiator`, and in particular they each encapsulate a geometric shape:

- The emitter shape corresponds to the volume of possible emission points. It can be either finite or infinite, if we are to model space scenarios, as explained in the induction paragraphs above. Currently, for finite shapes we only allow a single box, and an infinite emitter is necessarily the entire space; this has not proven to be a limitation so far. Beside its shape, the emitter also bears most of the physical information of the irradiation scenario: its directivity, e.g. unidirectional or isotropic, and what we call the "spectrum" or all the non-geometrical properties defining the emitted particles: a single energy and nucleus type for a monoenergetic emitter, or for instance energy-dependent flux files for a broad-spectrum emitter.

- The target geometry intuitively corresponds to the limits of "sensitive" silicon. By design choice, any generated ray that does not intersect the target is automatically discarded, therefore we impose the target shape to be finite in order to efficiently prune unnecessary rays. Its conceptual frontier can be a bit blurry: in SOI, the minimal sensitive volume to consider must enclose the silicon film of simulated MOSFETs; but if ions are modeled with a finite track radius, the target should bleed out a little, so as to not discard grazing ions that may deposit a non-negligible charge even though their track core does not intersect active silicon.

In bulk technology, since charge diffusion can happen on large scales in the wells or substrate where our transistors are sitting, the target must extend deeper and wider; it then has to be cropped somewhat arbitrarily for simulation times to remain reasonable. In both technology architectures, to alleviate the configuration phase for the end-user we propose an "autodetected" target box (once again, we use Manhattan-only geometry) satisfying the above criteria. Note that our target is necessarily a single box volume; for bulk technology this seems logical since the substrate and the embedded wells form a contiguous volume for carrier transport. In SOI, using multiple target volumes, e.g. one for each MOSFET, could be advantageous. However the idea underlying our emitter–receiver irradiation scheme is mainly to provide a first level of pruning; for example, rays falling in-between MOSFETs can be quickly discarded right after in the `Raytracer` module, incurring no significant computational overhead.

- The boundaries simply correspond to the limits of the simulated chip. As explained above, they are used to pull impact points back outside the simulation structure when using backward path-finding, before transporting ions from their emission point in the `Raytracer`. Note that this implicitly assumes that the chip is sitting in a vacuum—but if we were to model it in a surrounding medium such as air, we would just need to describe the associated geometry and material, and extend the "boundaries" accordingly.

Now as Figure 2.6 shows, the flow that we follow to construct one particle path is mainly determined by whether we run in forward or backward path-finding mode:

- When running in forward mode, the emitter is necessarily finite so that we can sample starting points in its volume. Outgoing directions are then sampled irrespective of the target, that is, the emitter is *a priori* oblivious to which region to aim for. Only then, after directions are sampled blindly, can we intersect the rays against the target volume and reject those that cannot reach the target. Forward irradiation is used for nuclear database scenarios (protons, neutrons), and alpha-particle simulation.

- In backward mode on the other hand, we pick an impact point and ingoing direction from the target, and as explained before the emission point is then deduced using the boundaries geometry. Backward irradiation thus models an infinite-emitter situation, where all constructed paths are kept. It is used for heavy ion scenarios, both unidirectional (beam testing) and isotropic (on-orbit situation).

Note that in both forward and backward modes, the spectral properties—i.e. the nucleus type and energy—are sampled at the end of the particle construction process, after the geometrical properties—i.e. the ray position and direction. This means we implicitly assume that there is no correlation between the geometrical and spectral variables. This choice would be a limitation if we were to simulate, for instance, both energy- and direction-dependent particle fluxes, which is sometimes taken into account in LEO where nearby Earth and its magnetic field can "cast a shadow" on some directions more than others—but these are very specific studies that generally fall outside our scope of silicon chip maker, and for all practical purpose we can consider the on-orbit spectra to be isotropic. Only nuclear database scenarios do bypass our logic of geometry-decorrelated spectrum, because we directly read both the energy and the direction of the particles in the list of nuclear reactions.

We may now consider Table 2.2 which describes the macroscopic properties of all our preset categories of irradiation scenarios. Each of them corresponds to a specific behavior of the sampling functions of the emitter and the target, and to a certain approach for the calculation of fluence Φ allowing the computation of consistent cross sections later on. Following (1.12) on page 31, Φ is generally defined as:

$$\Phi = \begin{cases} N_{\text{part}}/A_{\text{tgt}} & \text{in backward path-finding mode} \\ N_{\text{part}}/A_{\text{emi}} & \text{in forward path-finding mode} \end{cases} \tag{2.1}$$

Here N_{part} stands for the number of simulated particles, and A_{tgt} (resp. A_{emi}) is the *apparent* surface area of the target (resp. emitter) under the chosen directivity. What meaning we give to these terms in each particular scenario will also be explained below.

Irradiation scenario	**Path-finding mode**	**Emitter shape**	**Emitter directivity**	**Emitter spectrum**
On-orbit ions	Backward	Full space	Isotropic	Flux files
Backward ion beam	Backward	Full space	Unidirectional	Monoenergetic
Forward ion beam	Forward	Finite	Unidirectional	Monoenergetic
Alpha-like	Forward	Finite	Isotropic	Monoenergetic
Nuclear database	Forward	Finite	Reaction files	Reaction files

Table 2.2: Categories of supported irradiation scenarios.

2.3.2 The "on-orbit ions" scenario

Let us begin with the "on-orbit ions" scenario, being the test case that justified the implementation of a backward path-finding approach in the first place. As can be seen in Table 2.2, it corresponds to a situation where the particles' starting points can be anywhere in space, their directions are unrestricted—we further assume the distribution to be isotropic, and their spectral properties are dictated by flux files obtained from a CRÈME or SPENVIS simulation for instance. Under these assumptions, the distribution of impact points will be uniform on the target surface, as long as the target is a *convex* body: each point locally sees a fully-open sky above its surface, and thus all points receive the same amount of flux—with a non-convex object on the other hand, occluded points lying in cavities perceive a lesser flux because they are not irradiated by a full hemisphere of directions. In TIARA, the target we consider is an axis-aligned box, and sampling a point on such a surface is quite straightforward (see Section A.3.2, page 142). Hence the remaining question that arises is: what is the distribution of directions that the target perceives under an isotropic incoming flux? Or equivalently, how do we distribute chords in the target as if they were coming from an isotropic field of lines sent by an infinite emitter? As illustrated in Figure 2.7a, if we pick two points uniformly on the target's surface to generate a ray, we will obtain biased probabilities because the density of (uniform) lines joining two points is proportional to $\mathrm{d}\sigma\mathrm{d}\sigma' \cos\theta \cos\theta'/l^2$, favoring some directions more than others (the symbols stand for area differentials, normal angles, and chord length, as defined in the figure). This method only works for a spherical object as depicted in Figure 2.7b, where the chord length precisely cancels

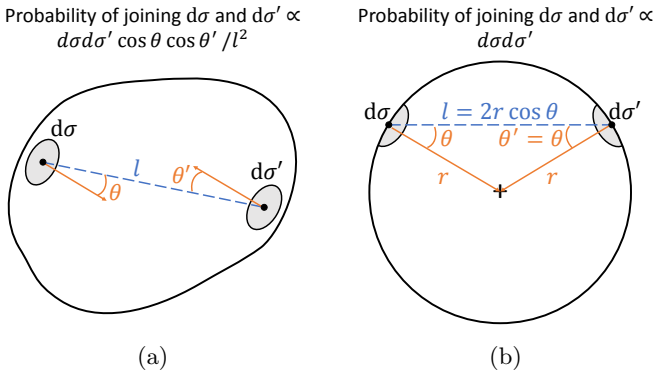

Figure 2.7: The problematic of sampling incident rays received from uniform lines in the isotropic irradiation case: (a) Joining two points sampled uniformly on the target generally does not yield the correct line density – (b) Unless the target is a spherical body. Adapted from [99].

out the cosine factors, giving a properly uniform line density only proportional to $\mathrm{d}\sigma\mathrm{d}\sigma'$. This can be the basis for generating rays incident on more complex geometries, by first sampling rays on a bounding sphere and then intersecting them with the enclosed shapes (see for instance [98]).

In TIARA however, as stated previously we consider an axis-aligned box for our target. From the geometrical considerations above, we see that the correct way to sample an ingoing direction after choosing the impact point, is to pick it in the local hemisphere at that point, with a probability density proportional to the *cosine* of the normal angle (the "local tilt"). This aspect is crucial: picking a uniform direction from the receiving point does not yield the correct density of lines, because uniform (isotropic) rays sent by an emitter are distributed with cosine weight from the receiver's standpoint. Intuitively, this is merely accounting for how much rays facing the surface are more probable than rays arriving at grazing angles. In terms of implementation, this is achieved by first generating a direction in the upper hemisphere with cosine-weighted probability (see Section A.4.3, page 145), then moving to a local coordinate system where the Z axis is the normal at the impact point, and the X, Y axes are constructed via a cross product with the normal vector. After that is performed, the ray's starting point on the circuit boundaries is easily deduced (Figure 2.8a), and all that remains is to sample the spectral variables for the particle. For an arbitrary distribution of nucleus types and energies, this is best done via the method of "inverse transform sampling" that makes use of cumulated densities (see Section A.2.2, page 140). Finally, in the formula for fluence computation (2.1) on page 59, the target effective area is its total surface divided by *four*. This apparent area once again comes from the cosine-weighted distribution of directions. More formally, the hemispherical integral:

$$\int \cos(\boldsymbol{u}, \boldsymbol{u_z}) \cdot \mathrm{d}\omega_{\boldsymbol{u}} = \int_0^{2\pi} \mathrm{d}\varphi \int_0^{\pi/2} \cos\theta \sin\theta \mathrm{d}\theta = \pi, \tag{2.2}$$

where $\boldsymbol{u}$ is the integration direction and $\boldsymbol{u_z}$ is the upward unit vector, is only a fourth of the complete sphere's solid angle 4π. In terms of implementation, we assign each target face with a weight w_i equal to its area A_i divided by four, and then add up these six weights to yield the overall apparent surface:

$$A_{\mathrm{tgt}} = \sum_{i=1}^{6} w_i = \sum_{i=1}^{6} A_i/4 \tag{2.3}$$

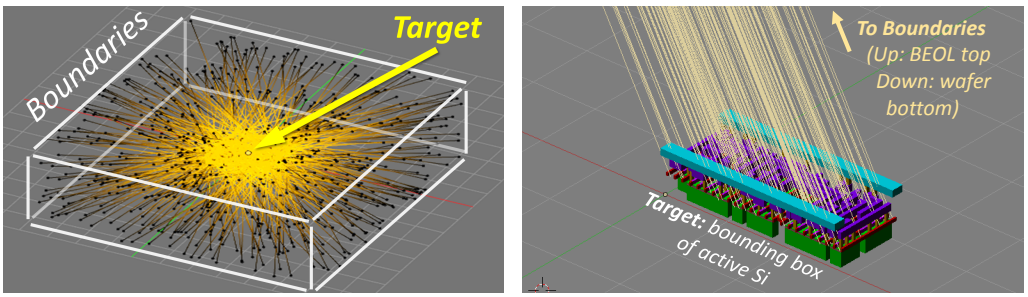

(a) (b)

Figure 2.8: Ion irradiation scenarios with backward path-finding: (a) Isotropic rays for on-orbit simulation – (b) Unidirectional rays for beam-testing simulation.

These are the same weights we use for the point sampling procedure mentioned above at the beginning of path construction (see the w_i coefficients in Section A.3.2, page 142). Proceeding this way also allows us to conveniently simulate scenarios with infinitely thick shielding, where rays can only come from above or below the circuit—the latter arising for a flip-chip situation for instance. Consider an upper-hemisphere irradiator for example: rather than using equal w_i weights to sample points on all faces and then having to *reject* paths pointing downwards, we simply multiply the upper face's weight by two (the sidewalls receive twice as less flux as the top) and the lower face's weight by zero (the bottom is fully shadowed), and sample impact points accordingly. Then when sampling directions, we flip the downward-pointing results and we *keep* the resulting paths, thereby yielding an acceptance rate of 100% without biasing the statistics.

2.3.3 The "ion beam" scenarios

Now that we have introduced all the formalism for our sampling procedures in the isotropic space irradiation scenario, the case of an ion beam simulated in backward mode is much simpler to consider (Table 2.2, page 59). Note that here the word "beam" only stands for the fact that we consider a unidirectional environment; what we are simulating in backward beam mode is not a local beam (which is performed by the "forward ion beam" scenario we discuss right after), but rather an unrestricted shower of parallel rays incident on the target (Figure 2.8b). Sampling a direction and an energy from a unidirectional, monoenergetic environment just means assigning a fixed, non-random value to those variables—or in other terms, sampling them from a Dirac distribution. The only non-trivial behavior is that of the surface sampling procedure: if the target is hit by a bundle of parallel rays, the distribution of impact points is not uniform as was the case for isotropic irradiation. It is, however, uniform on each target face separately, since the "local tilt" θ_i with respect to the irradiation direction is constant for face i. This is where our w_i face weights come at play: we set them to be proportional to the flux received by each face, i.e. with the cosine of the tilt (clamped to 0 for fully shadowed faces):

$$w_i = A_i \cdot \max(0, \cos\theta_i) \implies A_{\text{tgt}} = \sum_{i=1}^{6} A_i \cdot \max(0, \boldsymbol{n}_i \cdot \boldsymbol{u}_{\text{emi}}) \tag{2.4}$$

where $\boldsymbol{n}_i$ is the normal of face i, and $\boldsymbol{u}_{\text{emi}}$ is the emitter direction. Doing so, we produce the correct distribution of impacts, and we effortlessly compute A_{tgt} the target apparent surface perpendicular to the emitter direction, allowing to compute the fluence

from (2.1), page 59.

The next irradiation scenario in Table 2.2 (page 59) simulates a unidirectional, monoenergetic ion beam with forward path-finding. This means that unlike our previous infinite shower of parallel rays simulated backwards, we now consider a finite source emitting in a certain direction. The sampling procedures are very simple in this case: we sample the emitter volume uniformly for a starting point (see Section A.3.1, page 142), and then the direction and spectrum sampling methods just assign the associated variables to fixed values. Much like in backward beam mode, the fluence is then computed by dividing the number of emitted ions by A_{emi}, the emitter's transverse area w.r.t. the irradiation direction. Now because we are working in forward mode, the acceptance ratio for our produced rays is not necessarily 100%, unless we perform some trigonometry beforehand to make sure that the target is always reached. For arbitrary irradiation angles, one can easily see how impractical it could be to manually aim for the region of interest, when in fact this is precisely what the backward beam scenario achieves. Said differently, the forward ion beam scenario is superseded by the backward version, and in fact, this forward irradiation scenario was kept mostly for debugging purposes.

2.3.4 The "alpha-like" scenario

A more interesting use of forward path-finding is to simulate alpha-particle scenarios, or alike in terms of geometry (cf. Table 2.2, page 59): a finite radioactive source is sampled for emission points, and for isotropic behavior a direction is picked *uniformly* over the sphere of outgoing vectors (see Section A.4.1, page 143). Note the difference with backward path-finding, where we had to sample directions with cosine-weighted probability in the receiver's frame of reference. In the more "natural" forward approach of generating rays from emitter to receiver, the intuitive notion of isotropic irradiation is directly reflected in the implementation. The downside, of course, is that many rays may be lost because they do not reach the target[2]. Consider a large emitting slab of packaging material sitting on top of a 10-µm-thick BEOL stack covering a 0.1-μm^2 bit-cell. Even for emission points drawn closest to the target center, i.e. right above it, from simple solid angle considerations the probability of hitting the target is only about $0.1\ \mu\text{m}^2/4\pi(10\ \mu\text{m})^2 = 1/4000\pi$, which is less than 1 out of 10 000. When also considering emission points further away on the emitting slab, in TIARA typically we obtained acceptance ratios of 1 over 100 000 for alpha-particle scenarios using an isotropic source. A simple optimization is to use a hemispherical-directivity emitter to yield a speed-up of 2×. But a much more significant improvement was brought by implementing a "cone to target" mode where the outgoing direction is chosen uniformly inside a cone—or more rigorously, a spherical cap—directed towards the target (see Section A.4.2, page 143). The cone direction is thus dependent on the emission point at its apex, and therefore we have to use a fixed, global opening angle for all possible cones, for our rays to keep identical statistical weights—one of the major working hypothesis in TIARA, as explained in Section A.1.2 on page 138. This (half) opening angle is computed as:

$$\theta_{\max} = \arcsin\left(\frac{r_{\text{tgt}}}{\min d_{\text{etc}}}\right) \tag{2.5}$$

[2]These forward/backward considerations are also well known in the computational particle physics field, where Monte Carlo codes must simulate an arbitrary number of random scattering events for each particle trajectory. In that case, reverting the equations may become a lot more mathematically intricate than for our irradiator problematic, and adjoint, or reverse transport capabilities are often very valuable, sometimes even advertised in commercial tools.

where r_{tgt} is the target's radius, and $\min d_{\text{etc}}$ is the smallest emitter/target-center distance amongst possible emission points, computed by extremal considerations in Cartesian coordinates. In other words, at the closest emission point the cone of directions precisely encloses the target's bounding sphere (Figure 2.9a), thereby ensuring that the simulation is not biased by omitting potentially useful rays. To sum up, the "cone to target" directivity is nothing but an optimized implementation of isotropic irradiation, with a speed-up factor w.r.t. isotropic given by the ratio of the respective solid angles, i.e. $4\pi/2\pi(1-\cos\theta_{\max})$. This allows to reach success ratios of a few percents (Figure 2.9b), which is much more acceptable than the 10^{-5} order of magnitude given above for blind isotropic irradiation. This speed-up factor is accounted for by rescaling N_{part} accordingly in the fluence formula (2.1) on page 59, meaning that with a cone $1000\times$ narrower than the full sphere of directions, we are virtually shooting $1000\times$ more particles[3]. On a final note for alpha-particle scenarios, note that the emitter area we consider is that of its top or bottom face:

$$A_{\text{emi}} = X_{\text{emi}} \cdot Y_{\text{emi}}, \tag{2.6}$$

a reflection of our thin-slab assumption for such scenarios where particles almost never escape the emitter volume through the sidewalls. A more flexible implementation would be to let the user choose through which face to calculate the flux, or perhaps even to define a flux sensor of arbitrary location and orientation; for all other scenarios, the fluence can be defined unambiguously, but in the case of a finite, isotropic emitter, the flux follows an inverse-square law and the best we can do is propose preset configurations that make the most sense.

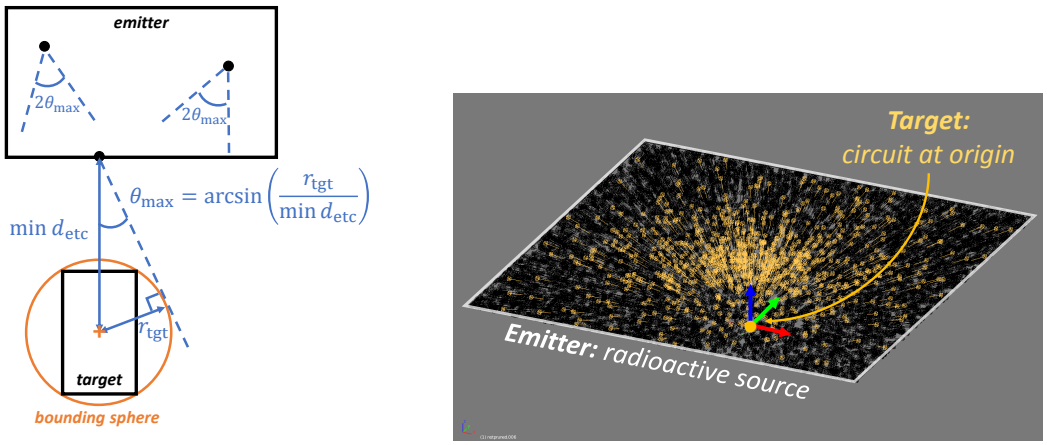

(a) (b)

Figure 2.9: Alpha-particle irradiation scenario with forward path-finding: (a) Schematic of the "cone to target" optimization – (b) Real-scale visualization of the accepted paths (orange) for a 60×60 µm^2 emitting slab on top of an 8 µm BEOL stack above the target.

[3]Note that this does not imply that alpha-particle simulations are made a thousand times faster by the "cone to target" speed-up; anticipating a little bit on parallelization by the `Driver` module discussed in Section 2.5.2, this just means that the `Irradiator` is not the limiting processing entity anymore. With blind isotropic directions, the `Irradiator` could be slower at generating valid rays than the `Raytracer` and the `Collector` were at processing them, thus "congesting" the simulation pipe and leading to slowdowns. Overall, the achieved improvement with the cone speed-up is in the $2-3\times$ range.

2.3.5 The "nuclear database" scenario

Last but not least, the "nuclear database" scenario (see entry in Table 2.2, page 59) constitutes our simulation scheme for neutron and high-energy proton environments. Using Geant4 particle transport simulations, nuclear reactions of protons and neutrons against silicon were recorded, i.e. by logging the direction, energy, and nucleus type of all ionizing products for each interaction event. An illustration is given in Figure 2.10 where we plot probability densities for geometrical and spectral properties of the simulated products, namely their outgoing tilt w.r.t. the primary particle's axis, and their initial LET in silicon. As can be seen, back-scattering is quite unlikely but does occur; it is necessarily the outcome of inelastic events, as elastic scattering can only result in angles below 90°. Then as far as ionizing power is concerned, there is typically a low-LET peak owed to light products such as protons and alphas, and a broader mode between 5 and 10 MeV$\cdot$cm^2/mg due to heavier nuclei such as the silicon recoils themselves. For a pure-silicon target material, the maximum LET that can be achieved is the Bragg peak of silicon, or 14 MeV$\cdot$cm^2/mg. The high-energy transport physics and simulation setup used were discussed at much greater length in [100]; here we will now try to give a notion of how these nuclear databases are incorporated in our irradiation flow, along with the necessary assumptions.

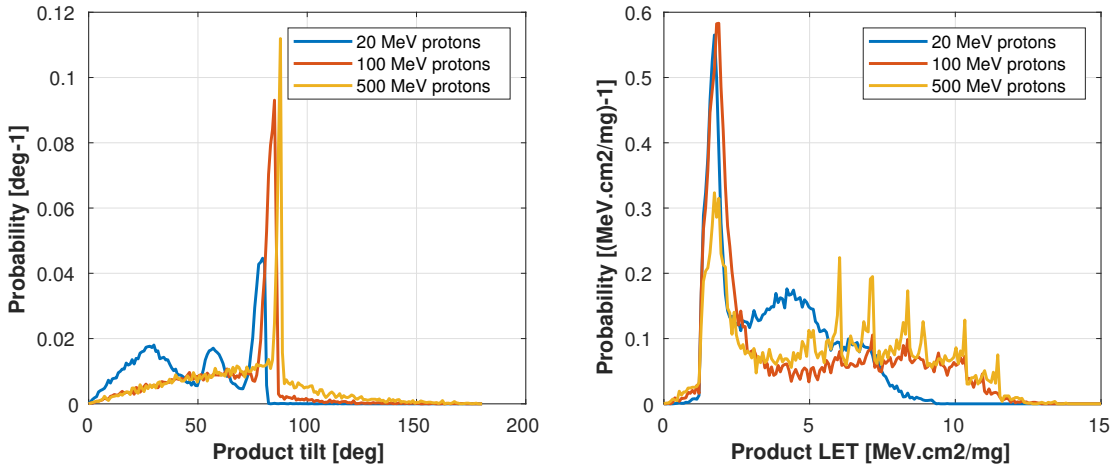

Figure 2.10: Statistics on several $p^+ - Si$ Geant4 nuclear databases. *NB: to avoid unequal weights due to different number of products across the reactions, tilt angle and LET distributions were computed for the most ionizing product of each reaction.*

The general idea in order to reuse Geant4 data in TIARA is to pick a random point for a nuclear reaction to occur, and then read a random entry in the table of reactions to generate ions. In order for this approach to work, the traveled medium has to be thin enough that cascading reactions can be neglected. This was ensured in Geant4 by targeting at normal incidence a silicon slab of small thickness ($Z_{G4} = 20$ µm) compared to the mean free path λ of the primary particles—recall from Section 1.2.1 on page 11 that the mean free path of neutrons is typically 20 cm, and so is that of protons for nuclear interaction. That being verified, we can rightfully assume that the resulting databases contains *independent* nuclear events to pick from in TIARA. To set the orders of magnitude, we typically generate Geant4 databases with a few hundred million primaries, leading to some dozens of thousands of reactions since $Z_{G4}/\lambda \sim 10\,000$—this could also be written in terms of nuclear cross section σ, remembering (1.3) on page 12. Because

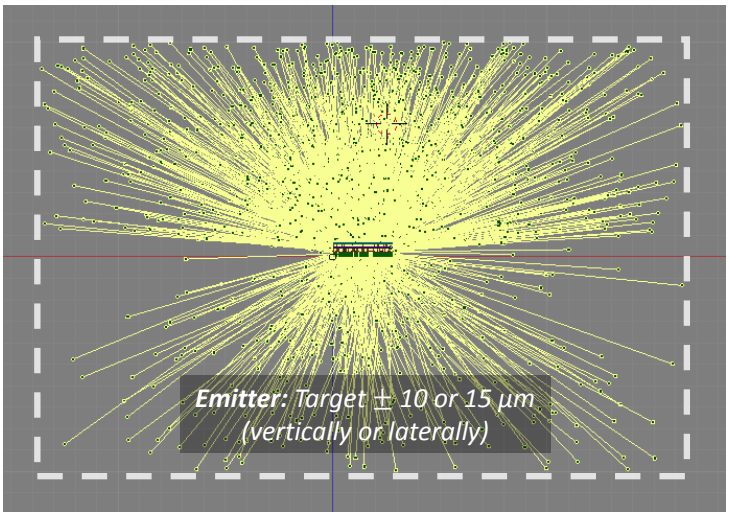

Figure 2.11: Nuclear database scenario (ionizing products *geometrically* reaching the target).

there are no multiple scattering events, we can also ignore altogether the location of the events issued by Geant4, and redraw them in TIARA with a *uniform density*—if we were looking at transport of the primaries on long distances with non-negligible attenuation, this would not be true.

Consequently, the sampling scheme in our nuclear database scenarios is as follows: draw a uniform location for a nuclear event inside the finite emitter volume, then a uniform integer indicating which reaction to simulate. After that, the emitter's direction and spectrum sampling methods are somewhat bypassed (see Figure 2.6 on page 57), in that they are only assigning to the generated products their directions and energies retrieved from the nuclear table, rather than drawing random properties from probabilistic distributions. In the mathematical sense, this means we are *resampling* from available data given by the nuclear files. Put differently, these databases provide us with *realizations* over the population of possible events, while in Section 2.3.2 the on-orbit flux files directly gave us the probability law to draw from. Note that the emitter volume we consider is a single box, whose dimensions are computed so that the autodetected target is enclosed with appropriate margins on all axes (Figure 2.11). By "'appropriate", we mean that once again a trade-off is to be found between having an excessively large emitter leading to slow convergence, and biasing the simulation by cropping the emitting volume too short. That being said, at present, there are three main limitations to our database approach for nuclear event simulation:

- Since the Geant4 simulations were performed at normal incidence, by reading the secondaries' directions in absolute coordinates, currently we only simulate vertically-incident protons or neutrons in TIARA. However, this limitation could be overcome by moving to a local coordinate system oriented in the primaries' frame of reference, or in other words rotating the secondaries' directions we read according to the primaries' incoming direction. Since the apparent thickness traversed by the primaries is then modified depending on their impinging angles, care has to be taken to avoid biasing the simulation. Still for now, the simulated proton or neutron fluence is computed from (2.1), page 59 by considering, firstly for the emitter area A_{emi}:

$$A_{\text{emi}} = X_{\text{emi}} \cdot Y_{\text{emi}}, \tag{2.7}$$

because our *emitter* volume that generates secondaries is itself a *target* to the normally-incident primaries—therefore for arbitrary directivity, the formula will be completely unified with (2.3) and (2.4) (pages 60 and 61) from backward path-finding scenarios. Secondly, the number of simulated (primary) particles $N_{\text{part}} = N_{\text{prim}}$ is given by a "Rule of Three" based on the number of simulated reactions N_{reac}, and the emitter thickness Z_{emi} compared to the mean free path λ:

$$N_{\text{prim}} = N_{\text{reac}} \cdot \frac{\lambda}{Z_{\text{emi}}} = N_{\text{reac}} \cdot \frac{N_{\text{prim,G4}}}{N_{\text{reac,G4}}} \cdot \frac{Z_{\text{G4}}}{Z_{\text{emi}}} \tag{2.8}$$

Again, for arbitrary directivity this will be modified to account for a potentially non-vertical traversed thickness. Note that combining (2.7) and (2.8), we could also derive the fluence of the primaries as the product of the reaction density (per unit volume) by the mean free path.

- Neutron databases were generated both for monoenergetic and white-spectrum neutrons, namely for JEDEC's NYC reference flux, LANSCE and TRIUMF accelerated spectra. Our proton databases on the other hand are all monoenergetic. As we will discuss in Section 4.3, this means that simulation results have to be convolved with the flux-versus-energy curve afterwards if we are to derive an SER. A possible evolution would be to generate proton databases for a few recurring environments, such as GEOs for a handful of longitudes assuming "regular" space weather. In a way, the problematic is very similar to that for accelerated beam testing, in that synthetic white sources are readily available for atmospheric neutrons but would not be very relevant for space proton spectra, the shape of which can greatly vary depending on the orbit.

- The target material used in our Geant4 simulations is pure silicon. Interactions with heavier elements such as metals in the interconnects can give rise to a broader LET spectrum, and in Figure 2.10 we would see a tail theoretically up to 34 MeV$\cdot$cm^2/mg (Bragg peak of copper) or even 89 MeV$\cdot$cm^2/mg (tungsten) [101]. For simulations on unhardened cells, this poses no issue as the SER is dominated by medium-LET contributions, as will be discussed in depth in Section 4.3. For robust cells this means that care has to be taken when analyzing simulation results, e.g. to avoid declaring that a certain design is proton-immune without having performed additional checks with heavy ions. Note also that interactions with lighter elements than silicon can be of importance as well. In [102], it was shown that front-face and backside irradiation results could sometimes differ because of the BEOL being mostly SiO_2 and not Si like the substrate. To sum up, while TIARA's results are not dramatically wronged by this Si-only assumption, it certainly constitutes our biggest shortcoming in treating nuclear interactions. We are currently generating Geant4 databases against other target elements to tackle this limitation. Note that this evolution will also require us to define multiple emitting volumes—or rather attach emission properties to our existing 3D boxes. To generate random emission points, the volume sampling procedure will then have to account for the statistical weight of each box, not only in terms of its volume, but also depending on the mean free path in the corresponding material—hence the volumetric density of reactions we mentioned earlier. Once the proper statistical weights are assigned to the emitting boxes, distributing points inside them is quite straightforward (i.e., merely a 3D adaptation of Section A.3.2 on page 142).

2.4 Radiation event processing

Following Monte Carlo generation of ionizing events by the `Irradiator`, the computation flow carries on with the `Raytracer` and `Collector` modules. These are the ones in charge of simulating the physical processes that will ultimately lead to injecting radiation currents in the `CircuitSolver` wrapper to the external SPICE simulator.

2.4.1 Ion transport in the structure: the `Raytracer` module

The function of the `Raytracer` is to compute charge deposition in the circuit by the ions received from the `Irradiator`. In terms of data structures, its inputs are an instance of the Circuit class, and Particle objects that carry the nucleus types, initial energies, starting points, and propagation directions. The task of the `Raytracer` is then to convert these Particle objects into "ParticleTrack" objects, or containers of segments that represent the ions' energy losses as they travel the chip structure.

These calculations are performed under a straight-ahead, Continuous Slowing Down Approximation (CSDA): we propagate the ionized nuclei along pure lines, and consider that their energy is transferred gradually at a continuous rate given by the total stopping power, rather than being the result of discrete collisions. Therefore, we solve (1.4) on page 13 for energy transfer by using tabulated values for the LET obtained from SRIM. In a way, this means that we consider the LET to be the cause and not the effect, when in the first place the stopping power is computed by averaging the energy losses simulated for a host of different trajectories. Mathematically, the range of a "virtual ion" experiencing the average energy transfer is not necessarily the average range of ions experiencing individual scattering events. But apart from low energies dominated by nuclear stopping, the CSDA range yields a very close approximation to the average range. This allows us to decouple the geometrical aspects from the physical mechanisms of energy transfer, and leads to a two-pass algorithm implemented in the `Raytracer` (see Figure 2.12a):

- First, we trace rays against the 3D structure for geometrical intersections where materials could change. This is done via a J5geolib routine testing for line–box intersection, which we call on all boxes in the 3D model. The result of these operations is a collection of segments each belonging to a single box, representing the possible ion tracks if the nuclei were to propagate unattenuated across the chip. Note that currently, our data structure for the 3D model is not designed as hierarchical—although some boxes are regrouped by well or implant to help `Collector` calculations later on. The boxes are stored in a flat array and therefore, when looking for intersections the traversal time is linear with the total box count. This could be revised by indexing the circuit geometry in a Bounding Volume Hierarchy (BVH) data structure instead, then leading to logarithmic complexity for intersection tests. To date, this has proved unnecessary, mainly because we simulate infinitely thin ion tracks represented by a single line. Thus ray tracing calculations have never become the limiting factor so far, even when simulating structures of several dozen thousand boxes; however as we will discuss in Appendix B, using a BVH may become necessary if we want to simulate tracks of finite radius, carried by many more segments.
- We then solve the energy transfer equation (1.4) via forward Euler method: we move along the potential segments created before in small Δl steps, resulting in

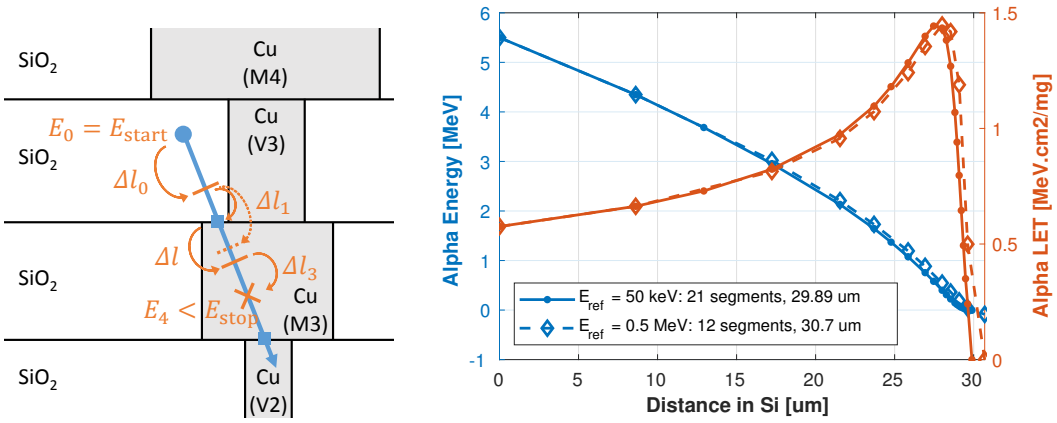

(a) (b)

Figure 2.12: Working principles of the `Raytracer` module: (a) Splitting of the ion path at geometrical intersections (blue) with potential material interfaces, and further refinement of the segments to track the energy losses (orange) – (b) Finite-difference solving of the energy transfer equation for a 5.5 MeV α-ray in silicon, with different settings for the adaptive step-size control.

energy decrements $\Delta E = E - E'$:

$$\Delta E = LET(E) \cdot \Delta l \tag{2.9}$$

where the SRIM stopping power is in unit energy per unit length, and is taken at E, or the energy value *before* the step (explicit Euler scheme). To control the accuracy of the integration, the grid step is computed to avoid excessive energy loss between two points:

$$\text{Find } \Delta l \text{ such that } \frac{\Delta E}{E + E_{\text{ref}}} \leq 10^{-d} \tag{2.10}$$

where E_{ref} stands for a certain reference energy, and d is an accuracy in digits, although non-necessarily an integer. Therefore, when the energy is large ($E \gg E_{\text{ref}}$), Δl is sized based on a *relative* energy loss criterion ($\Delta E/E \leq 10^{-d}$); and at small energies where dividing by E would be numerically unstable, the criterion becomes an *absolute* criterion ($\Delta E \leq 10^{-d} E_{\text{ref}} = \Delta E_{\text{abs}}$). Note that the Δl obtained from (2.10) is in fact a trial value that we confront to the remaining distance until the next intersection—the actual step we take being the minimum of both. For these transport calculations we use the *total* stopping power, but we also record its *electronic* part in our internal representation for the track segments: it is the deposited charge, or the energy lost to δ-rays, that is relevant for the `Collector` to operate on. Finally, we also define a cut-off energy E_{stop} below which we consider that the ion has slowed down to rest, and ray tracing calculations can end for this nucleus[4].

[4]In Figure 1.7, page 14, it can be observed that at low energies, the LET follows a power law in the energy: $\partial E/\partial x \approx -KE^{\alpha}$. If the α exponent was greater than 1, integrating the differential equation would lead to an asymptotic decay with E never reaching zero, and our cut-off energy E_{stop} would be a necessity. But from Figure 1.7 we have $\alpha \approx 1/2$, meaning that even under the CSDA—which is quite questionable when collisional losses dominate—we will find that the energy does reach true zero. Therefore E_{stop} (typically 1 eV) is merely a commodity to avoid out-of-range accesses on SRIM tables.

Other than the SRIM tables, the `Raytracer` module does not require much configuration: the geometrical ray tracing pass does not rely on any free parameter. As for the physical pass for the energy losses, configuration is mostly a one-off task which consists of finding an optimum in terms of accuracy versus the number of generated segments, by tuning the d "digits" parameter and E_{ref}. For d we use a global setting of 1.0 typically (split the segments around 10% energy variation), and E_{ref} has to be sized for each nucleus type. This is exemplified in Figure 2.12b, where a too small E_{ref} results in excessive refinement near stopping (the absolute variation criterion is dropped), while a too large E_{ref} would lead to creating too many segments at high energies (the relative variation criterion would be overpowered). Thus, for a given ion the optimal value for E_{ref} is typically a fraction of its Bragg peak, in the transition regime between electronic and nuclear stopping. On a final note regarding our adaptive stepping algorithm: the Euler method being a first-order procedure, if the LET is constant the differential equation for energy transfer is solved in a single step with no error, and in general the error is quadratic in the step size. This means that instead of basing our stepping strategy on the energy variation along segments, we could very well consider the LET variations instead, resulting in larger segments of near-constant LET rather than near-constant energy. Although this would reduce the number of generated segments, this would relax the refinements essentially at high energies, where the Δl is rather geometry-limited anyway; geometrical intersections set a limit to the achievable gain in total segment count, and thus our energy-controlled segments seem satisfying enough at present.

Figure 2.13 illustrates the variety of nuclei and energy ranges that the `Raytracer` may encounter. The color map gives the LET depending on the ion's atomic number and on its kinetic energy per nucleon, in MeV per amu. Superimposed is a point cloud (in red) for typical recoils on a LEO (p^+–Si products at 100 MeV). It exhibits the two-hump distribution of light and heavy fission fragments, the latter being often generated

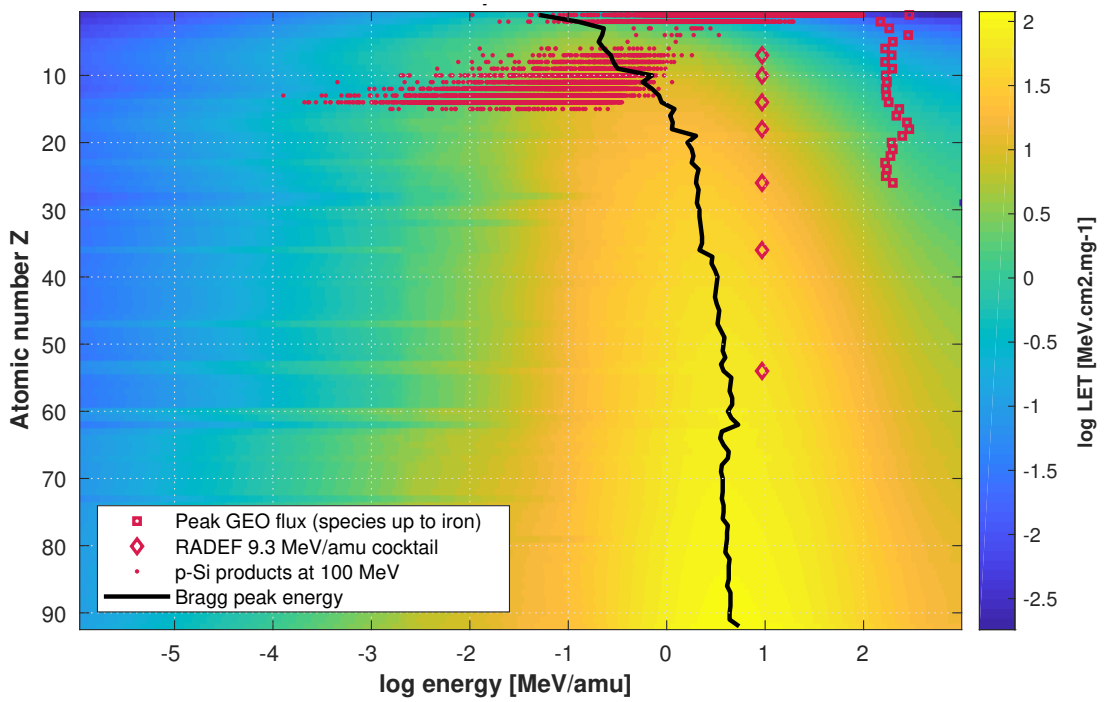

Figure 2.13: LET color map for all ion species in silicon (SRIM data used by TIARA's `Raytracer` module). Bragg peak energy versus atomic number Z drawn as a continuous line for clarity.

with sub-peak energies—therefore they are short-ranged particles that justify the need to adequately ray trace their trajectories against a precise 3D model of the BEOL and FEOL. At the other end of the spectrum lie the heavy ions encountered in GEO (red squares), for which we only indicate the energy of maximum flux by ion species, for the sake of clarity. When not stopped by shielding, they travel the sensitive silicon volumes with close to no attenuation at all. But note that their emulations in cyclotrons have much lower energy (see the 9.3 MeV/amu cocktail at RADEF in red diamonds), causing quantifiable differences in ion track structure, in spite of having LETs covering that of GCRs. Put another way, cosmic rays are hard to reproduce "in the lab", even with the best particle accelerators—we will elaborate on the implications of this in Appendix B.

2.4.2 Charge transport and circuit response: the `Collector` and `CircuitSolver` modules

Following charge deposition calculations in the `Raytracer`, the `Collector` simulates the carrier transport mechanisms leading to charge collection at electrical contacts, and how this affects the circuit once the `CircuitSolver` injects radiation currents in the SPICE netlist.

In terms of APIs, the `Collector` receives "ParticleTrack" objects from the `Raytracer`, and it operates by delivering "RadiationSource" instances describing the time-dependent currents induced on the affected devices. These sources can either be attached to p-n junctions, or to MOSFETs directly (recall Figure 2.5 on page 55 for the circuit internal representation). As illustrated in Figure 2.14a, the former case corresponds to bulk technology where transistor implants form a vertical junction with the substrate or well underneath, while the latter represents our modeling choice in FD-SOI where electron–hole transport happens mostly in the lateral direction, along the source–channel–drain axis. For both transistor architectures, the presence of these time-dependent radiation sources often means that we are essentially *shunting* the impacted devices momentarily: the high concentrations of free carriers turn an off-state FET into somewhat of a conducting wire, even though this is not controlled by the transistor gate as in normal operation.

As was explained in Section 1.4.2 on page 35, custom models are required in order to capture numerically-challenging semiconductor physics whilst meeting severe speed constraints in the context of Monte Carlo SEE simulation. The `Collector` implements several models to support different technology architectures and speed–accuracy trade-offs. From a programming standpoint, each model corresponds to a specialized class inherited from an abstract Collector class. In planar bulk technology, we use the "Diffusion–Collection" model introduced in [7] and used throughout [5]. As briefly explained in Section 1.4.2, the model uses an analytical solution to ambipolar diffusion in neutral regions (a spherical Gaussian Green function) to derive the carrier density near the collecting junctions; then assuming that the current density is only due to electric drift at the edge of the depletion region, the current is obtained by integrating the current density over the contact area. The Diffusion–Collection model is thus very computationally efficient, typically requiring less than one second for evaluation, which is not bottleneck compared to SPICE simulation. Beside this, integration of a Random-Walk Drift–Diffusion (RWDD) model, whose concept was presented in [103], is currently ongoing in TIARA. The model simultaneously simulates drift, via a deterministic velocity imparted by the electric field, and diffusion with a random-walk algorithm or Brownian motion. The main benefit of RWDD compared to the Diffusion–Collection model is its

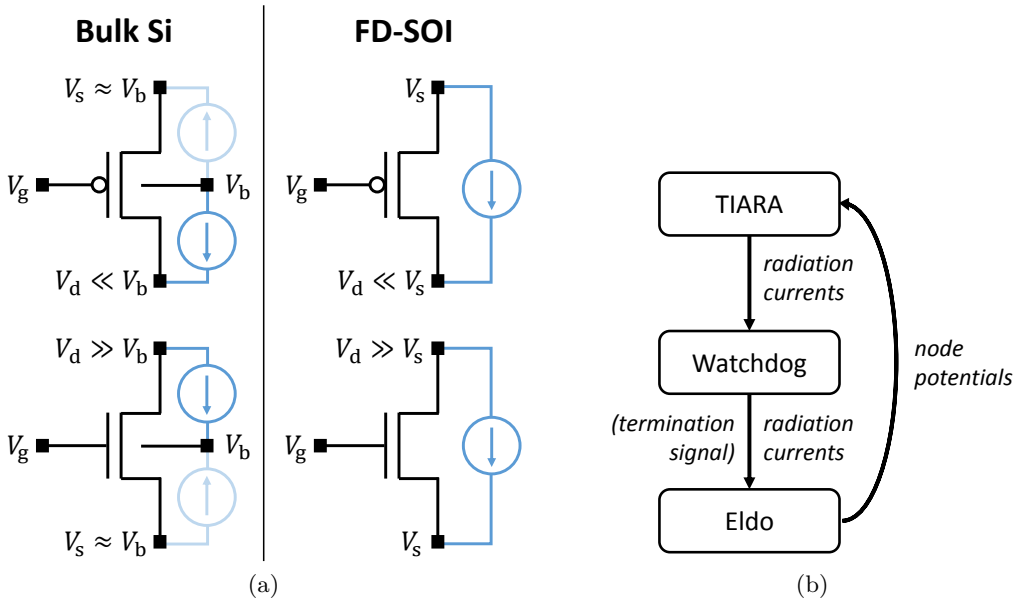

Figure 2.14: Main tasks fulfilled by the `Collector` and `CircuitSolver`: (a) Schematic view of the computed radiation currents, flowing mostly between drain implant (n^+, p^+) and bulk electrode (n-well, p-well or p-substrate) in bulk Si technology, and between drain and source in FD-SOI. *NB: current source symbol always oriented in the positive current direction* – (b) Fully-coupled simulation flow, with watchdog process in-between TIARA and Eldo running in interactive mode.

true circuit coupling, owed to the fact that the junction fields are modulated depending on the circuit node potentials. This allows to capture many of the desired dynamic properties of the radiation–circuit–radiation feedback loop, as was highlighted in [103]. Then concerning FD-SOI technology, two original models were developed within the framework of this thesis, and we will come back to them in Chapter 3. The commonality of all these algorithms is that they are *compact* SEE models, and therefore the physical inputs used by the `Collector` are oftentimes "effective" parameters to describe the semiconductor physics at play—such as an overall diffusivity for collection models in bulk. These parameters are either extracted from TCAD simulations, or some are taken as degrees of freedom when calibrating the models directly against silicon data. An exception is the advanced FD-SOI model that we develop in Chapter 3, which directly operates on the doping profiles of the MOSFETs and not uniform, global dopant concentrations as is the case for other models. The doping profiles can either be obtained from a TCAD process simulation, or with construction analysis methods such as Scanning Capacitance Microscopy (SCM) or Secondary-Ion Mass Spectroscopy (SIMS).

Next, the task of the `CircuitSolver` is to inject these radiation stimuli in the SPICE netlist of the circuit. The main dependence is thus a Process Design Kit (PDK) containing, among others, the MOSFET models needed for the simulation to run. Wrapped inside the `CircuitSolver` is in fact Eldo (Mentor Graphics) performing the SPICE simulations [8], and thanks to special compatibility modes we can also run HSPICE (Synopsys) netlists and PDKs [35]. When we run in sequential mode, radiation currents are fully computed beforehand based on the circuit state right before injection. Then the radiation waveforms are inserted as PieceWise Linear (PWL) current sources, and the `CircuitSolver` uses a single system call to run a transient Eldo simulation from start to end. On the other hand, with circuit-coupled collection models we use Eldo in interactive mode, a mode which allows to run a simulation step by step, with dynamic

access to simulation parameters at all times. The radiation currents are then refreshed based on the evolving node potentials, and the other way around. To enable this, the TIARA process is pipelined with the Eldo process and *vice-versa*, meaning that the `stdout` or standard output stream of one process is redirected via a Unix "pipe" to the `stdin` input of the other. Note that to avoid a zombie Eldo process in case the TIARA process crashes, actually we take care in adding a watchdog process in the pipeline, as illustrated in Figure 2.14b. The watchdog is in charge of passing the radiation stimuli from TIARA onto Eldo, and also to kill the Eldo process in case TIARA terminates unexpectedly.

Overall, Eldo's interactive mode is an extremely powerful feature, in that it allows the user to interface SPICE with virtually anything. This offers perspectives for simulation at higher abstraction level on large blocks (i.e. the same way mixed-signal simulation works), or as is our case here, for lower-level dynamics requiring physical accuracy—we could very well code an entire mixed-mode TCAD simulator with this pipelining ability. This is a related, but not completely similar approach to the Verilog-A language (the continuous-time subset of Verilog-AMS [104]) for custom analog circuit modeling: we compute Δf variations on certain quantities over small Δt time steps, which means we need to do the job of discretizing our equations and the TIARA–Eldo feedback is not native; Verilog-A allows tight coupling by offering direct access to the differential operators manipulated by the circuit solver, and therefore compact models can be described with differential equations as relationships on $\partial f/\partial t$, leaving the solver in charge of the discretization procedures to yield a numerical solution. The drawback of Verilog-A is that models are meant to be rather compact, thus the language is not so suited to handling large amounts of data. For the model we develop in Section 3.3 for example, large arrays will be needed to solve semiconductor equations on a grid.

2.5 Result aggregation and simulation harness

In addition to all the simulation modules that have been described up to this point, two non-physics modules are in charge of outputting intelligible results and controlling the execution flow: the `Analyzer` module performs post-processing on the event simulations and tally functions in order to derive statistical quantities such as the cross section or the error rate; the `Driver`, or simulation harness, orchestrates proper execution of the simulation chain with high-level calls to the individual submodules, both during standalone and parallel execution.

2.5.1 Per-event extractions and macroscopic quantities: the `Analyzer` module

The `Analyzer` is in charge of three main tasks: first of all, to perform the necessary treatments on the output waveforms from electrical simulations, so as to detect SEUs and SETs and measure their individual characteristics. Then, it manages an event database in order to allow the user to inspect correlations of as many event-defining input and output parameters as possible. Finally, it computes characteristic statistical metrics like the cross section, based on the simulated fluence or flux and probes at user-defined circuit nodes.

At the end of every impact simulation, SEU and SET detection is performed at user-defined probes, by comparing the simulated waveforms on the defined nodes to "golden"

signals obtained from a reference simulation without any injected current: if the final voltage differs significantly from its expected value—via a configurable threshold—then an SEU flag is raised. If at any time, the voltage signal swings by more than a certain amount, and for more than a certain duration, compared to the golden run, then an SET flag is issued, and the pulse amplitude and width are also measured. We measure pulse width both with respect to an absolute amplitude (e.g. the amount of time where the swing magnitude is more than $V_{\mathrm{dd}}/2$) and with respect to the glitch's own amplitude V_{SET} (e.g. the amount of time where the swing magnitude is more than $V_{\mathrm{SET}}/2$, i.e. the FWHM).

This post-processed data allows to populate an "EventDatabase" object, which will be used later on by the tally functions to compute σ or the SER. Moreover, in order to keep track of as many parameters as possible, this EventDatabase and the associated APIs are in fact visible of all TIARA modules, allowing them to record useful fields on a per-event basis. Indeed, all intermediate simulation files cannot be kept on disk: circuit simulation outputs can be very large, especially if many nodes are probed in post-layout simulations with parasitic RC networks. Mere SPICE logs can be heavy as well, especially when the netlist is elaborated by flattening all circuit hierarchy, including sometimes the entire PDK. Therefore, the general philosophy in TIARA is to perform automatic removal of event simulation files provided the user wishes so, but to configure all modules such that at least the minimum defining set of inputs–outputs is recorded for each event: one should be able to always trace the effects back to their causes. In practice, this is implemented by "EventResult" containers whose size varies as the processing chain goes, where the `Irradiator` records for instance impact locations, the `Raytracer` may save the overall deposited charge, to which the `Collector` may add a collected charge, and finally the `Analyzer` measures an SET width—although in the last case, this could conceptually be a task for the `CircuitSolver` to perform. The EventDatabase can then be dumped to a CSV file at simulation end, for custom inspection and correlations. Note that future evolutions could include the use of an SQL database, which would make queries more flexible. To date however, this has proved unnecessary.

This EventDatabase is also used to compute the statistical SEE metrics: σ is computed exactly as in (1.11), page 31, for as many SEU or SET probes as necessary, and based on the fluence returned by the `Irradiator`. Then the SER is computed as:

$$SER = \sigma \cdot \phi_{\mathrm{ref}} \tag{2.11}$$

where ϕ_{ref} is reference flux held by the `Irradiator`. This calculation is the direct reflection of (1.13), page 31, in that the convolution product is already performed by directly simulating all the energy and orientation content of the environment spectrum. In other words, the Monte Carlo procedure performs the multivariate $\sigma * \phi$ integral on its own, and all the `Irradiator` needs to know is how much real time has been simulated.

2.5.2 High-level control flow: the `Driver` module

Having described how all calculations are sequenced in the simulator, it would be unnecessary to retrace all the linear simulation chain at this stage. In fact, running one TIARA impact in standalone mode, i.e. on a single workstation, pretty much corresponds to successive calls to all modules, in the order of the above paragraphs. Distributed execution, on the other hand, ought to be described.

For parallel computing, TIARA is articulated around a master–slave architecture, in

which the "master" sends batches of particles to "slaves" for processing. Communications are done via files written and read on shared network disks. Note that more elaborate solutions, such as using sockets on network ports for communications, have been deemed too adherent to a particular infrastructure. So far, the disk exchanges have not proved to be a serious limitation, although latencies can sometimes be observed. In practice, TIARA's binary can either run the `Driver` in master mode, or slave mode (and of course standalone mode). The master performs all upstream operations needed to create the circuit (i.e. with calls to the `Builder3D` and the `LVSmatcher`), whose description is then saved on disk. The master then runs the reference (golden) simulation for analyses later on. The `Irradiator` is also unique, and part of the master process, so as to avoid handling multiple pseudo-random sequences: when running on a computer cluster with dynamic load balancing and possible crashes, this keeps the simulation as deterministic as possible and ensures maximum reproducibility of the results. Thus, the `Driver` requests particle batches to its `Irradiator`, which will be sent to the slaves.

Upon first particle batches, slaves are created by submitting TIARA slave jobs (`Driver` in slave mode) to LSF. Eldo license check-in is performed upon job submission, via a special resource request. Further ahead, this will mean that the SPICE process managed by the `CircuitSolver` is simply forked, and not sent to a different physical machine. Slaves load the circuit description from disk, and then the event processing sequence *per se* can begin for all slaves and batches: slaves loop over the `Raytracer`–`Collector`–`CircuitSolver`–`Analyzer` sequence, and write the post-processed data to event result files. Put another way, slaves are responsible for all the per-event treatments, including the associated *micro-analyses*. On the other hand, when one batch is fully processed, the master's `Analyzer` reads results from the associated slave, updates the event database, and computes (updates) the required *macroscopic* aggregates. The simulation carries on until the total requested even count has been reached.

As a technical consideration, note that batches have a controllable size of a few hundreds or thousands of particles, which results from two constraints: if batches are too large, the simulation may not reach a proper "steady state" with constant processing throughput between the different nodes; for instance, with a batch size close, but not equal to the total number of desired events divided by the number of slaves, one slave may be assigned two batches when all others only have to process one, and the simulation will have to last for about twice longer. Small-sized batches, on the other hand, lead to a better refresh rate for the outputs. However, since they are currently dispatched to the slave nodes as text files, with overly-small batches, disk accesses may slow down the simulation overall.

With this parallelized architecture, TIARA has proved to work robustly with a handful to hundreds of slave nodes, yielding SER simulation times as low as a few minutes for simple circuits under simple environments. As was mentioned previously, LSF constitutes one of our most specific dependencies (together with Eldo). Thus it is difficult to assess how TIARA can perform on other infrastructures, other than by practical deployment, and this certainly leaves room for future opportunities.

2.6 Comparison with the state of the art

Having given an in-depth description of all the major aspects that make up for TIARA's present-day code, a detailed comparison with other Monte Carlo SER codes is now presented to further elaborate upon what was introduced in Section 1.4.2, page 35. A sum-

mary of our literature review is given in Table 2.3, where we have tried to fill as many of the blanks as possible, but sometimes leaving interrogation marks to avoid speculative "best-guesses". Note that beside TIARA from ST, two other foundry-developed tools should be mentioned, namely IBM's Soft Error Monte Carlo Model (SEMM) version 2 [105, 106], and Intel Radiation Tool (IRT) [107]. These two simulators were not included in the comparison however; for the former, to the best of the author's knowledge all accounts of the code have been published prior to the 2010s, thus making timely comparison quite delicate; for the latter, the lack of published data other than [107] makes it hard to populate the comparison table with generic, test-case independent features. Note however that the authors are among the few to tackle FinFET technology, and a major highlight of their work is the proposal of a bias-dependent (or "circuit-coupled", in this manuscript) SEE model for bulk FinFET, implemented in Verilog-A and derived from original ideas in [68].

Starting with the physics aspects, it can be noted that none of the codes implement both a high-energy transport code *and* a semiconductor device simulator, except Particle and Heavy Ion Transport System (PHITS) coupled to HyENEXSS, developed by JAEA [108]. As already stated several times, this has to do with the fact that brute-force TCAD is the bottleneck in the simulation chain, compared to high-energy transport. In [108] the authors have to use aggressive mesh optimizations and impact pruning (i.e. not running the simulation when the deposited charge is below a certain cut-off), in order to yield a tractable simulation of about a thousand impacts. They report a simulation on a single device (an nMOS to represent the most sensitive device in an SRAM), which suggests that the method could suffer scale limitations. On the other hand, in MRED [3, 65] the nested sensitive volumes, modeling a space-dependent collection efficiency, yield nearly instantaneous results as to the semiconductor response, and simulations have been undertaken on large circuits such as a DICE [61]. However, no circuit feedback on collection has been implemented in MRED, to the best of our knowledge.

Circuit feedback is also one shortcoming in MC-ORACLE [4,63], MUSCA SEP3 [63, 64], and TIARA's `Collector` for bulk technology: the Diffusion–Collection model [7] and its refinements can by essence not implement true circuit coupling, since the analytical diffusion solution assumes zero external electric field (only the effective collection velocity could be made to vary w.r.t. voltage). Note that these codes, and TIARA's current version as well, are based on static nuclear databases for neutron and proton interaction in a limited set of targets, oftentimes silicon only. This certainly constitutes a limitation that needs to be tackled. In [109], radial ion track profiles and longitudinal energy loss straggling (i.e. LET fluctuations around average) were implemented in MUSCA SEP3, overcoming limitations of the simple CSDA transport procedure, and thus capturing many desired properties of detailed ion track structures obtained in Geant4. This is a direction we are currently taking in TIARA developments, and more light will be shed on these topics in Appendix B.

Now with the circuit-coupled collection models, apart from PHITS-HyENEXSS already examined, two commercial tools feature an interesting approach. Accuro [23, 110] from Robust Chip Inc. proposes an optimized mixed-mode 3D device simulation—more specific information as to the simplification procedures are not given however; the computational burden of TCAD is then reported to be alleviated by one to two decades, making SER estimation very intensive, but tractable. Then, TFIT [111] from IROC Technologies derives response models in collaboration with foundries, by running vast TCAD explorations across impact coordinates, LETs, and device biases. For each supported

technology, this produces a large database of transient currents to interpolate upon. Circuit coupling is *emulated* by effective modeling of the current plateau features [112,113], i.e. the fact that charge collection is clamped once the voltage has dropped following the initial injection. Although in TIARA, our general philosophy is rather to keep transport models quite compact (and the original FD-SOI models developed afterwards fall in this category), these two approaches prove to be very effective at solving the semiconductor problem. A peculiarity in the approach of Accuro, is the fact that only constant-LET ions are considered, thus neglecting all high-energy transport phenomena. This is conceptually similar to the original set of approximations in the RPP model [60], the benefit being that once the cross section is computed across LETs, positions, and orientations, SER projections in arbitrary environments is not just "fast", it is instantaneous.

Leaving aside the physics, the importance of software aspects cannot be overstated. Proposing a GUI, as commercial tools Accuro and TFIT do, is of course a definite plus; beside this, something we must insist on is that flow automation is key to an accelerated workflow. This is definitely a direction taken by TIARA along the course of this PhD. With automated 3D structure creation based on the circuit layout, and the electrical-to-physical correspondence enabled through our custom LVS scan, at constant CPU resources the simulation throughput truly is multiplied. Overall, "user-friendliness" can be just as important as physical accuracy, and tools such as MRED or MC-ORACLE proposing a web interface, immediately gain popularity—and once an enthusiastic community builds up around a tool, soon additional developments follow. In that respect, perhaps the number one limitation of TIARA could be its proprietary character—although external deployment is not at all excluded on the mid-term.

	MRED [Weller2003, Warren2009, Reed2013]	**MC-ORACLE** [Correas2008, Wrobel2011]	**MUSCA-SEP3** [Hubert2009, Artola2011, Raine2012]	**PHITS-HyENEXSS** [Abe2012]	**Accuro** [Lilja2016, www.robustchip.com]	**TFIT** [US20080077376A1, private communications]	**TIARA 2011** [Uznanski 2009, Autran 2012]	**TIARA 2018** [this work]
Institution(s)	Vanderbilt University, NASA/GSFC	Université Montpellier 2	ONERA	JAEA	Robust Chip Inc.	IROC Technologies	STMicroelectronics, IM2NP	STMicroelectronics, IM2NP
Accessibility	Publicly available	Academic licensed tool	?	?	Commercially licensed tool	Commercially licensed tool	Proprietary internal tool	Proprietary internal tool
Nuclear interactions	Geant4-coupled, any projectile or target	DHORIN database (n in Si/SiO_2)	Geant4 database (n, p^+ in Si)	PHITS-coupled, any projectile or target	Database approach (n in Si?)	Database approach (n in Si?)	Geant4 database (n in Si)	Geant4 database (n, p^+ in Si)
Direct ionization	Geant4-coupled, detailed ion track structure and FEOL/BEOL reported	CSDA transport with SRIM LET, preset BEOL/FEOL stack	CSDA transport with SRIM LET, track structure emulation. Preset BEOL/FEOL stack	CSDA transport with PHITS LET, preset BEOL/FEOL stack	Constant-LET ions for fast SER convolution	Initial LET with SRIM + range cut-off, preset BEOL/FEOL stack	Initial LET with SRIM + range cut-off, preset BEOL/FEOL stack	CSDA transport with SRIM LET, track emulation in SOI. Layout + process-stack based BEOL/FEOL.
Carrier transport, circuit response	RPPs with Q_{crit} or double-exponential in SPICE. No circuit feedback on collection	RPPs with Q_{crit} or SPICE with Diffusion-Collection. No circuit feedback on collection	SPICE with Diffusion-Collection. No circuit feedback on collection	Optimized full 3D TCAD simulation. Native circuit feedback on collection	Optimized mixed-mode 3D device sim. Native circuit feedback on collection	TCAD-database interpolated current in SPICE. Circuit coupling emulated (plateau)	Diffusion-Collection / double-exponential (funneling) in SPICE. No circuit coupling	+ empirical FD-SOI double-exp. in SPICE (no circuit coupling) or circuit-coupled FD-SOI 1D device simulation
Supported circuits (assuming CMOS)	Generic circuit (Q_{crit}, RPP, waveform configuration). SRAM array, DICE reported	SRAM arrays reported	SRAM arrays, inverter chains reported	Single device reported	No layout/netlist restriction. Hundreds of simultaneous devices reported	No layout/netlist restriction	Templated 3D coordinates/netlists for SRAM arrays & DFFs	No layout/netlist restriction. Hundreds of simultaneous devices simulated
Supported technologies	Q_{crit}, RPP, waveform adjustments reported in Bulk/SOI/FinFET	Planar bulk	Planar bulk	No restrictions (HyENEXSS configuration)	Planar bulk, preliminary FinFET reported	TSMC 40G, 28HP, GF40LP	Planar bulk	Planar bulk, FD-SOI
Supported environments	No restrictions (Geant-4 configuration)	Terrestrial (vertical n, isotropic α) and unidirectional HI	Terrestrial (vertical n, isotropic α) and cosmic rays (isotropic HI, unidirectional p^+)	No restrictions (PHITS configuration)	Direction- and LET-dependent flux	Terrestrial (vertical n, isotropic α) and unidirectional HI	Terrestrial (vertical n, isotropic α) and unidirectional HI	Terrestrial (vertical n, isotropic α) and cosmic rays (isotropic HI, unidirectional p^+)
Typical time per impact (assuming full processing)	< 1 s with RPP, seconds with SPICE	< 1 s with RPP, seconds with SPICE	Seconds	?	Minutes (10x-100x faster than general-purpose TCAD)	Seconds	Seconds	Seconds
Setup automation, flow integration	Manual RPP setup. Popular CRÈME-MC online interface	Batch mode and online interface	Manual inputs for physical/electrical structures	?	GUI available	GUI available, automated physical/electrical mapping (LVS)	Semi-automated inputs for phys./elec. structures, parallel computing (LSF)	Automated 3D struct., automated phys./elec. mapping (LVS), parallel comp. (LSF)

Table 2.3: Comparison of Monte Carlo SER simulation codes.

Chapter 3

Transistor level modeling of single-event effects

This chapter addresses the work carried out at the lowest layer of abstraction, i.e. the radiation response of elementary transistors. One of the primary objectives in this research work being to understand the 28 nm FD-SOI technology from a single-event perspective, we first discuss TCAD studies in FD-SOI in Section 3.1. This allows us to delineate what sets FD-SOI apart from bulk silicon or Partially-Depleted Silicon On Insulator (PD-SOI) technologies. We then move onto the contributions in terms of modeling work carried out during this PhD program. In other words, based on the physical insight obtained from TCAD *simulation*, we develop compact single-event *models* for use in TIARA, which hopefully retain the necessary physics whilst meeting our speed requirements for Monte Carlo soft-error simulation: Section 3.2 presents *Bipolar Amplification Basics (BA-Ba)*, a lightweight behavioral model that was developed to account for the main features of FD-SOI radiation-wise. Then Section 3.3 goes through the derivation of a more elaborate model called *Carrier Transport with Humble One-Dimensional Equations (CaTHODE)*, which captures these mechanisms from first principles and allows to extend TIARA's scope of applicability to more circuits and operating conditions, as will be discussed in Chapter 4. Both BA-Ba and CaTHODE are new `Collector` classes inside TIARA, the latter especially representing one of the main innovations brought within the framework of this thesis.

Note that throughout this chapter, comparison with experimental results is mostly indirect: as stated in Chapter 1, unless with dedicated test structures and off-chip apparatus such as external oscilloscopes with very high bandwidth, at advanced technology nodes it is difficult to obtain measurements of transistor-level effects like transient current or voltage waveforms. In the industrial context of this thesis, oftentimes the goal has been to characterize the radiation sensitivity of either standard or robust cell libraries, using test vehicles designed to measure cell-level SER. Therefore, experimental comparison will have to come in the next chapter, when we show heavy-ion and proton cross section results on sequential and combinational logic. Thus in the present chapter, we consider TCAD—though not an absolute truth—our lowest-level, most accurate reference for investigating physical mechanisms and parametric trends.

3.1 TCAD characterization of 28 nm FD-SOI

In order to comprehend the charge transport mechanisms pertaining to 28 nm FD-SOI technology, 3D TCAD simulations were performed using the Sentaurus software suite from Synopsys [31], and published in [10]. The present section is an adaptation of the TCAD study from that paper. The two main axes we wish to investigate in the technology are first of all, the sensitive volume for charge collection, and second of all, the bipolar amplification phenomenon.

3.1.1 Study objectives, previous work, simulation setup

The effective sensitive volume is known to be very *thin* in 28 nm FD-SOI thanks to the UTBB dimensions, as already mentioned in Section 1.5.1 (page 39) and associated publications. What needs to be investigated, is how it extends on the *other dimensions*. It has long been known [114] that the most radiation-sensitive regions are the reverse-biased junctions of a transistor, where the electric field is strong enough to separate the electron-hole pairs and induce a parasitic current. In planar bulk technology, these areas are primarily located under the drain junction of the off-state transistors. In 28 nm FD-SOI however, because the BOX lies immediately beneath the drain, we expect the sensitive region to be a lot *narrower*.

The other crucial parameter in SOI is bipolar amplification, or the activation of the parasitic lateral Bipolar Junction Transistor (BJT) by the incoming radiation [12]. In an nMOS transistor for instance, as illustrated in Figure 3.1, when electron–hole pairs are deposited by an ionizing particle, they get separated by the electric field of the junctions; while the electrons are quickly collected at the drain, the holes are repelled inside the body and can accumulate in neutral regions, therefore raising the body potential and triggering the parasitic source–channel–drain BJT. As a result, the source starts injecting electrons that give rise to a drain current, and this phenomenon lasts for as long as the holes linger in the channel: because they are majority carriers, they do not recombine quickly. This results in a parasitic bipolar gain—which for SEE considerations is usually defined by the ratio of the collected to the deposited charge—above one.

In PD-SOI technologies, strong parasitic bipolar gains (up to 100×) can arise because the deep BOX leaves room for large neutral regions and thus majority carriers accumulation [116, 117]; this can occlude the benefit of having reduced sensitive volumes compared to bulk technology. Local body ties can be employed to mitigate the floating-body effects (known as *kink effects* in the analog world), at the cost of reduced integration density [118]. For fully-depleted SOI, previous work down to 50 nm has reported better behavior than PD-SOI in terms of bipolar amplification [119,120]. However at the beginning of this PhD work, no published data existed in 28 nm FD-SOI, to the best of the author's knowledge. Most recently, a detailed study of the collection mechanisms was published [121] on the same technology as presented here, showing amplification factors consistent with the trends highlighted in the present work. Note that much earlier FD-SOI studies, not specific to radiation, had also been conducted in [115], but the silicon film thicknesses involved at the time were not representative of modern technologies. The parasitic bipolar gains reported were thus much larger than present day technologies, although the authors did demonstrate that film thickness was a primary driver in the overall response.

To quantify all these effects, we use 3D TCAD structures of elementary transistors, which are brought to an off-state steady state with a drain bias and no gate bias, and

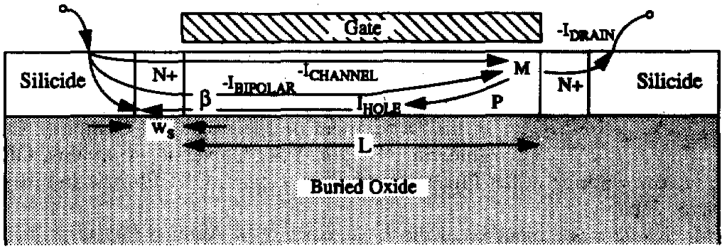

(a)

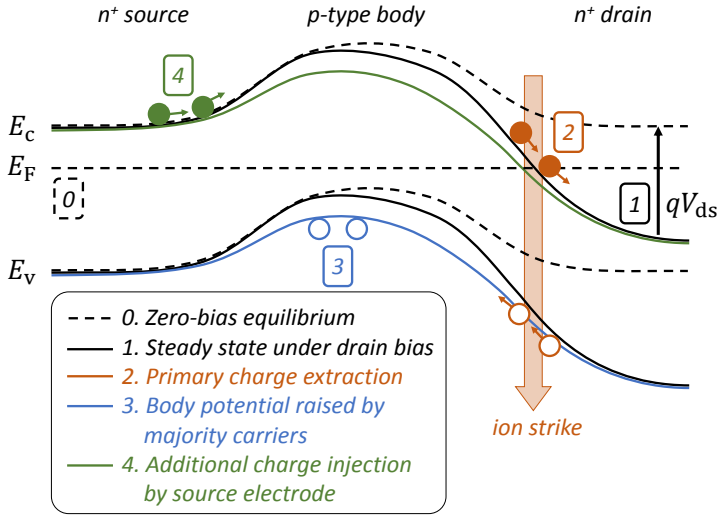

(b)

Figure 3.1: Parasitic bipolar amplification in FD-SOI (a) illustrated in physical space (taken from [115]) and (b) drawn in a band diagram.

on which numerous (vertical) impact simulations are run (Figure 3.2) for ions of various LETs. The ion tracks are modeled with a carrier generation term G_{ion} of Gaussian dependence in space (R_{ion} = 10–100 nm depending on the ion) and time (temporal extent T_{ion} = 2 ps):

$$G_{\text{ion}}(\boldsymbol{r},t) \propto LET \cdot \exp\left[-\left(\frac{t-t_{\text{ion}}}{T_{\text{ion}}}\right)^2\right] \cdot \exp\left[-\left(\frac{||(\boldsymbol{r}-\boldsymbol{r}_{\text{ion}})\times \boldsymbol{u}_{\text{ion}}||}{R_{\text{ion}}}\right)^2\right] \tag{3.1}$$

where $\boldsymbol{r}_{\text{ion}}$ is the starting point of the ion track, $\boldsymbol{u}_{\text{ion}}$ is the unit vector defining its direction, and the proportionality factor is a normalization constant making G_{ion} dimensionally consistent with a number of pairs created per unit time, per unit volume.

Our structures are calibrated for hydrodynamic transport. Outside the radiation term, in our simulation setup the main activated models are:

- *Band structure, carrier statistics:* We assume Fermi–Dirac carrier statistics, and BandGap Narrowing (BGN) arising at high doping densities is accounted for via the Del Alamo model [122]. This locally shrinks the energy gap into:

$$E_{\text{g,eff}} = E_{\text{g}} - qV_{\text{bgn}} \tag{3.2}$$

with straightforward notation. Accounting for BGN is of paramount importance for proper modeling of the parasitic BJT: under Maxwell–Boltzmann statistics for the

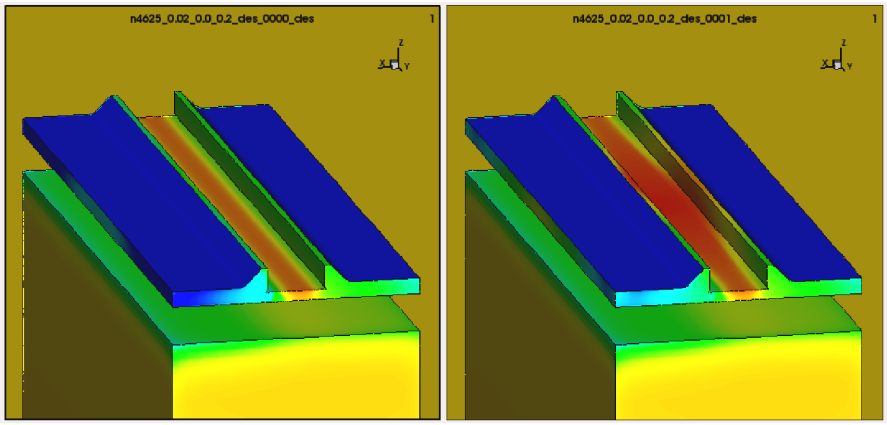

Figure 3.2: Simulated hole density before and during an ion strike on a 28 nm FD-SOI nMOS.

sake of presentation, when the equilibrium mass-action law is modified into:

$$(np)_{\mathrm{eq}} = n_{\mathrm{i,eff}}^2 = n_{\mathrm{i}}^2 \exp(qV_{\mathrm{bgn}}/k_{\mathrm{B}}T) \tag{3.3}$$

(where n, p are the electron and hole densities, n_{i} is the intrinsic carrier concentration, k_{B} is Boltzmann's constant, and T the absolute temperature), then with highly-doped implants the minority concentration is raised by an exponential factor. Thus, overlooking BGN would yield overly pessimistic projections for the parasitic bipolar gain, and hence the total collected charge, under radiation events.

- *Mobility:* We use the Philips unified mobility model proposed by Klaassen [123] for mobility degradation with temperature, doping and, more importantly, carrier density, i.e. to account for Carrier-Carrier Scattering (CCS). Indeed, CCS has been widely recognized as one of the driving mechanisms for radiation-induced carrier transport given the very high injection levels at play [114], and is known to be a major source of uncertainty in ion track modeling [124, 125].

- *Recombination:* We consider Shockley–Read–Hall (SRH), i.e. impurity-assisted recombination depending on the doping level, and Auger i.e. direct band-to-band recombination depending on the temperature. Auger recombination especially sets in at high LET, and properly captures recombination in the track core.

- *Generation:* We use the Van Overstraeten model [126] to account for impact ionization, or avalanche multiplication of the carriers under high electric fields. This mechanism is especially important for a proper rendition of bipolar effects in FD-SOI, since it can assist the amplification and lead to larger multiplication factors.

3.1.2 Sensitive volumes extraction

The first step in investigating the sensitivity of FD-SOI is to find the areas that collect radiation-induced charge the most. To quantify this, we perform impact simulations with positions swiping across the transistor. We extract the drain current $I_{\mathrm{d}}(t)$ and integrate it in time to measure the collected charge:

$$Q_{\mathrm{coll}} = \int_t (I_{\mathrm{d}}(t) - I_{\mathrm{off}})\mathrm{d}t \tag{3.4}$$

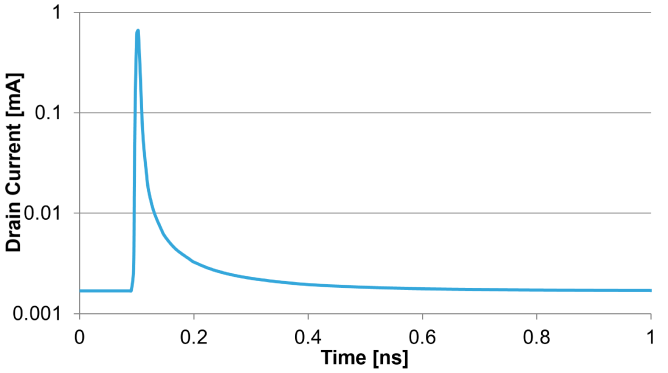

(a)

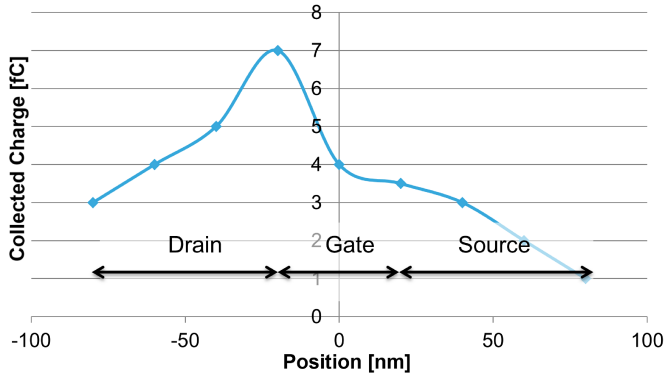

(b)

Figure 3.3: FD-SOI sensitive volumes extraction: (a) Simulated current for an nMOS impact – (b) Collected charge versus impact position for nMOS.

where we must subtract leakage current I_{off} in order to avoid a diverging integral, and to truly isolate the radiation collection from the transistor's background operation. Figure 3.3a displays an example of current transient we obtain. Note how fast the collection process is, lasting essentially 20 ps (the y-scale is logarithmic). Figure 3.3b summarizes the collected charge data obtained versus impact position. As was to be expected given the fact that the (good-conducting, neutral) implants host no significant electric field, the most sensitive area is located at the drain–channel junction. This is where the *lateral* electric field is largest, unlike in bulk silicon technology where the entire drain area is subject to strong collection due to the *vertical* junction field. In a nutshell, not only is the sensitive volume in FD-SOI very thin thanks to the ultra-thin active film, but it is also very narrow.

3.1.3 Characterization of bipolar amplification

Now, to quantify the parasitic amplification in 28 nm FD-SOI, the bipolar gain β is calculated as the ratio between the collected charge Q_{coll} defined in (3.4), and the deposited charge Q_{dep} defined as the integral of the ion track density over the SOI volume:

$$Q_{\text{dep}} = \iint\limits_{t,\ \boldsymbol{r}\in\text{SOI}} G_{\text{ion}}(\boldsymbol{r},t)\cdot \mathrm{d}^3\boldsymbol{r}\mathrm{d}t \tag{3.5}$$

$$\beta = \frac{Q_{\text{coll}}}{Q_{\text{dep}}} \tag{3.6}$$

Note that this is just one possible definition for the so-called "bipolar gain": we could also define the bipolar gain as in [127], by taking a ratio of currents. Our integral definition here has the advantage of giving a simple figure that *relates* to bipolar amplification, but amplification cannot be entirely isolated from the transistor's normal behavior, which is why we need to subtract the leakage current I_{off} in integrating for the collected charge. Also, impact ionization assists the injection of carriers by multiplying them in high-field regions, and so the β ratio we define here should be understood as an overall metric encompassing many transport and collection effects all at once, rather than the actual bipolar gain of a bipolar junction transistor. In order to distinguish the collected carriers that come from injection at the source (possibly multiplied by avalanche effects) from those that directly originate from the ion track, one would have to follow the individual carrier trajectories and scattering processes in a Monte Carlo solution of the Boltzmann Transport Equation (BTE). But in a fluid description such as with the hydrodynamic model used here (or with a drift–diffusion solution), such a distinction is impossible since we are solving for continuous carrier densities.

Figure 3.4 exhibits the β data obtained for nMOS and pMOS impacts at the drain–channel junction, under different back biasing conditions (the body potential is affected by biasing the substrate beneath the BOX). Outside very low injection, the effect of body bias is moderate. For LETs above 1 MeV$\cdot$cm^2/mg, or typical values for the SEU threshold in memory cells, as will be discussed later on, the parasitic gain is bounded by 3, and in the very high LET regime it drops to 1: under very strong injection conditions, the large, but finite current injected by the parasitic BJT becomes negligible before the amount of deposited charge. Thus the majority of charge collection comes from the initial injection, and because this massive charge entirely neutralizes the electric field, no strong avalanche multiplication occurs, which explains why overall the β ratio becomes one at high LETs. These low values for the parasitic bipolar gain in FD-SOI are in good agreement with other work in the literature (see for instance [128]). All in all, the parasitic amplification is very well contained in 28 nm FD-SOI when compared to counterpart PD-SOI technologies which frequently exhibit β ratios of 10 and more [128]. The reason for this is that the full depletion avoids majority carriers from accumulating in neutral body regions, as they are quickly expelled by the electric field of the space charge region. Note that β essentially depends on the 2D cross section of the transistors (gate length, film thickness...). Furthermore, within one technology platform with "fixed" gate length, variations in transistor sizes are mainly allowed in the "width" dimension by extrusion. This means that our simulations on elementary transistors remain valid for various layouts as long as the process is unmodified.

3.2 Behavioral single-event model for FD-SOI: BA-Ba

The TCAD study carried out above has allowed to underscore two main properties of 28 nm FD-SOI: first of all, thin and narrow sensitive volumes have been delineated, showing that the maximum collection "yield" lies at the drain–channel junction where the lateral electric field is most intense. Secondly, the parasitic bipolar amplification is quite small in that technology, and its variations have been computed over a wide LET range. To capture these two main specificities, and enable SER projections within a certain validity range that we discuss hereafter, a simple compact model was developed for radiation-induced currents in FD-SOI, called Bipolar Amplification Basics (BA-Ba). Note that this section adapts elements from [13].

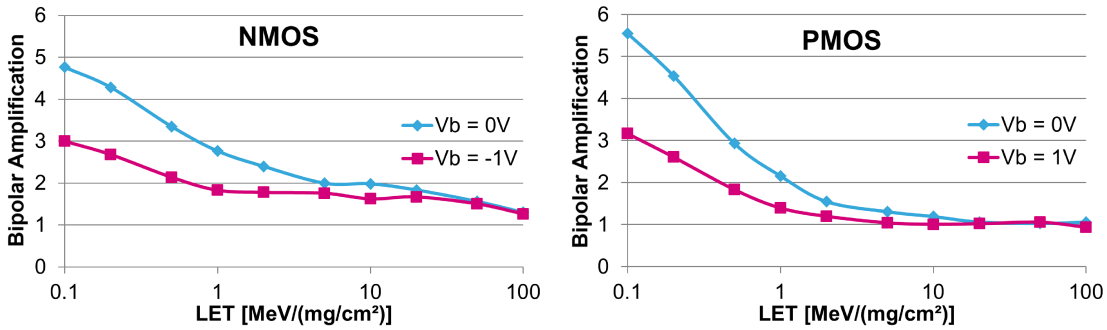

(a) (b)

Figure 3.4: Bipolar amplification for nMOS (a) and pMOS (b) 28 nm FD-SOI transistors.

3.2.1 Constitutive equations and modeling choices

In [10], it was shown that by applying a uniform β ratio to the charge deposited in active silicon and using a critical charge criterion derived from SPICE simulations, satisfying α-SER predictions could be obtained in TIARA. Alpha particles only cover a narrow LET range, however. Expanding on this work, we devise our response model BA-Ba with an LET-dependent bipolar amplification $\beta(LET)$, modulated by a spatially-varying collection efficiency $\eta(x)$ along the length axis of each sensitive transistor. To capture the non-uniformities of this charge collection yield, TIARA's `Raytracer` module must discretize the ion track with nanometric resolution typically, at points $\boldsymbol{r}_i = (x_i, y_i, z_i)$ where the x coordinate is along the FETs' length axis. The total collected charge is then computed as:

$$Q_{\text{coll}} = \sum_i Q_{\text{dep}}(\boldsymbol{r}_i) \cdot \beta(LET(\boldsymbol{r}_i)) \cdot \eta(x_i) \tag{3.7}$$

In other words, the collected charge is retrieved by assuming bold linearity for the sake of simplicity. Using this method, for any impact we are able to calculate Q_{coll}, which is then related to a double-exponential current waveform of rise and fall times τ_{r} and τ_{f} by:

$$I(t) = \frac{Q_{\text{coll}}}{\tau_{\text{f}} - \tau_{\text{r}}} \cdot \Big[\exp(-t/\tau_{\text{f}}) - \exp(-t/\tau_{\text{r}})\Big] \tag{3.8}$$

and this is what ultimately feeds the SPICE solver.

As a summary, the "free" parameters in our response model are the LET-dependent bipolar amplification, the position-dependent yield, and the current waveform rise and fall times. To fix all of these, the TCAD simulations presented previously provide us with orders of magnitude and relative trends, and for fine tuning, as will be presented in the next chapter (Section 4.1.1), we rely on calibration against heavy-ion data at normal incidence on SRAM cells. In practice, this means that:

- Initial values for $\beta(LET)$ are taken from the bipolar amplification data shown in Figure 3.4 (page 85), and then adjusted to match experimental cross-section measurements, especially w.r.t. threshold LET.

- Likewise, initial values for $\eta(x)$ are obtained from the curve in Figure 3.3b (page 83), after normalization to obtain a weight between zero and one: $\eta(x)$ it set to 1 at the drain–channel junction, and goes to 0 further away. Note that although the drain is strongly doped, charge collection does not vanish immediately inside it. Calibration

of $\eta(x)$ is mainly driven by the measured saturation cross sections. Note that for other gate lengths than $L_{\min}$ shown in the previous TCAD results, the function is simply stretched by the $L_g/L_{\min}$ ratio. This was confirmed by TCAD simulations, showing only small variations in collection for other transistor lengths.

- τ_r and τ_f, the rise and fall times of the injected double-exponential waveform, are both below 10 ps, representative of the current shapes computed in TCAD. Since the critical charge, roughly speaking, scales as $C.V + I.\tau$, modulating such small values of rise and fall times with the LET or the position would make little difference to the end result.

Note that radiation currents obtained with BA-Ba will be plotted afterwards in Section 3.3, when we compare our two compact SEE models with TCAD references.

3.2.2 Validity range

Immediately following the implementation and calibration of this response model, the question that arises is that of its validity range. Firstly, of direct relevance for circuit designers, is the accuracy of the modeled charge collection across PVTs. In BA-Ba's current implementation, no process dependence (i.e. slow nMOS, fast pMOS obtained by corner-case doping to act on the mobility) or temperature variation is encoded, and this certainly leaves room for improvement. Had it not been for the more elaborate collection model presented afterwards, which can natively handle these parameters, additional developments in BA-Ba would have been mandatory. As for the 'V' in PVT, further TCAD simulations suggest that the supply voltage V_{dd} has weak influence on all of the above parameters (e.g. 10% decrease in Q_{coll} for -50% V_{dd}): while in bulk technology, the depletion depth expands with V_{dd}, resulting in deeper collection, in depleted technologies like FD-SOI the collected charge is primarily dictated by the—rigid—sensitive volume. In the current implementation of BA-Ba, the injected currents are not made dependent on the impacted FET's initial V_{ds}, apart from a cut-off value (0.1 V) below which we consider that no collection occurs. Note that these trends were recently corroborated by [121], showing that around $V_{nom} = 0.9\ V$, the overall charge amplification factor is quite flat with respect to drain–source voltage, down to at least 0.6 V. Then at voltages below 0.3 V, the authors report a sharper drop in collected charge. Thus for all practical purpose, our step-function V_{ds} dependence is well justified, unless we are to target ultra-low voltage applications where the arbitrary cut-off voltage could be challenged.

A more intrinsic limitation of BA-Ba is that the model is not circuit-coupled: the double-exponential current may depend on the impacted FET's V_{ds} *at injection time*, but even then a frozen waveform is injected in the netlist, and no further circuit feedback on charge collection can occur. This means that BA-Ba cannot reproduce the so-called "plateau" effect [112,113], wherein charge collection is clamped once the separating electric field extinguishes due to the voltage drop on the collecting contacts. Nevertheless, non-coupled collection models have been widely used to predict SEU rates, and likewise, later on BA-Ba will be shown to perform very well for SEU prediction in memories and sequential cells. This is because in two regimes, SEU decision is not dramatically made wrong by omitting the circuit feedback on charge collection: when largely below the SEU threshold, the injected current does not lead to a significant voltage drop, which means that in retrospect, computing and injecting a bias-independent current was quite justified. Conversely, if the LET is far above LET_{th}, when using a bias-independent current source, the predicted voltage can become largely unphysical, but the fact remains

that a significant voltage drop does occur, and hence an SEU is issued: shutting off the current source as the voltage difference vanishes, would make no difference to the computed SER. Now, in the "gray zone" close to SEU threshold, the decoupled assumption leads to pessimistic predictions for SEU occurrence, if the injected waveform precisely matches the current that would be collected by a hard-biased device. However, this can usually be compensated by free parameters in the model, so as to correctly match the SEU/no-SEU frontier. For SET simulation on the other hand, using bias-independent currents yields pessimistic predictions in terms of pulse width, since the current does not vanish after the voltage has dropped.

Therefore, as a circuit-decoupled collection model, BA-Ba is mainly entitled to SEU simulation. Its main benefit, compared to the model presented right after, is its speed: computing Q_{coll} and $I(t)$ from (3.7) and (3.8), page 85, is nearly instantaneous compared to the SPICE simulation, or in other words, charge collection calculations are never the bottleneck in the simulation chain.

3.3 Advanced single-event model for FD-SOI: CaTHODE

In this section, we now introduce the second single-event model that was developed for FD-SOI during this PhD, called Carrier Transport with Humble One-Dimensional Equations (CaTHODE). Briefly speaking, CaTHODE is based on a custom 1D drift–diffusion solver dynamically coupled to SPICE, with TIARA's `Collector` operating Eldo in interactive mode. This present section is an adaptation of [14].

3.3.1 Previous work, motivations

Within an ideal Monte Carlo SER simulation chain, all physical processes ought to be accounted for with the highest degree of accuracy permitted by state-of-the-art tools. As complex as it may be to determine the cross sections for nuclear interactions and various scattering processes, once said cross sections are well characterized, the problem of energetic particle transport is amenable to pragmatic resolution via simple forward approaches. While not necessarily the most statistically optimal, this means that millions to billions of radiation events can routinely be simulated, with properly clusterized computers [65]. The semiconductor problem of radiation-induced charge transport and collection, on the contrary, is inherently very stiff, because the constituting equations of electrostatics and carrier motion are tightly coupled. The formulated system solved by TCAD tools is computationally very intensive, typically requiring several hours for a single impact simulation on a 3D structure. Thus in this "ideal" Monte Carlo SER simulator we described, TCAD is bound to be the bottleneck element. It is neither suitable for error rate prediction nor interactive feedback at circuit-design time: both these applications require single-event models compact enough to execute in seconds, yet physically accurate. Within the fuzzy notion of "physical accuracy", two peculiar features ought to be tackled by such models if they are to be used at analog-circuit ("SPICE") level: time dependence, and circuit coupling.

Circuit-level simulation of single-event effects in FD-SOI has been performed as early as the 1990s in works such as [129] for instance. The Ambipolar Diffusion with Cut-off (ADC) model [130] offered extremely valuable insight into the collection mechanisms, but the analysis was carried out under steady-state conditions only. That same idea of considering purely ambipolar transport in the absence of electric field, followed by collec-

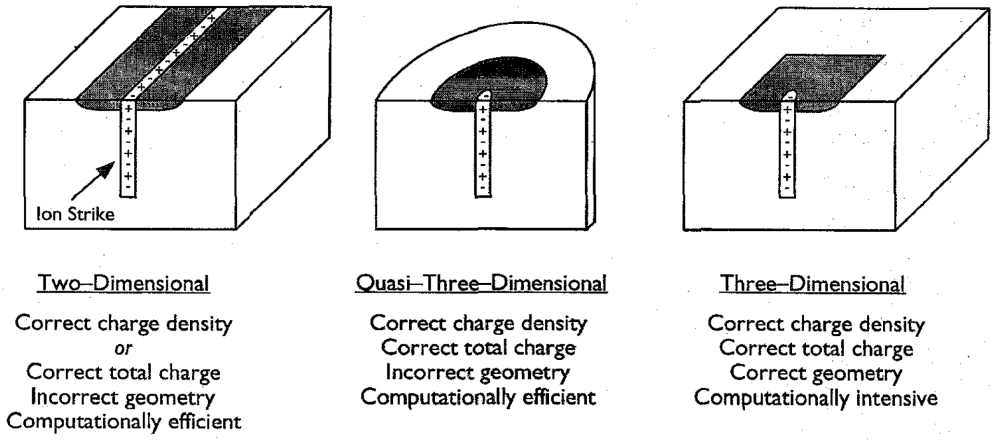

Figure 3.5: The dimensionality curse for TCAD simulation *in bulk technology*. Taken from [114].

tion from the electric field in a depletion region, gave rise to the "Diffusion–Collection" model [7], which is used in TIARA for bulk technology: the time-dependent diffusion equation can be solved analytically, however due to the separation between diffusive and drift transport, no true circuit coupling can be implemented. The authors of [68] were among the first to tackle the question of circuit coupling, replacing the (independent) current source by a bias-dependent current-generating network comprised of a capacitor and a (dependent) current source. Lumped circuit-element formulations have also proved to be popular for SOI technology, where parasitic bipolar amplification is commonly accounted for by injecting the primary radiation-induced current into the base of a BJT [131, 132]. Equivalent-circuit approaches like these can offer great predictive power, but only after proper calibration, which is perhaps their main weakness when compared to more physics-based models.

The single-event model developed in this section is based on direct resolution of the semiconductor equations on a one-dimensional mesh. The thrust behind this "1D" assumption is twofold: firstly, it alleviates computational cost, meeting the speed requirement for SPICE-level single-event models in TIARA. Secondly, it is physically relevant to modern fully-depleted technologies such as FD-SOI or FinFET, which both exhibit active silicon volumes whose dimensionality can be reduced: unlike bulk technology where the transport problem is intrinsically three-dimensional, in 28 nm FD-SOI the "thickness axis" can actually be collapsed thanks to the ultra-thin (7 nm) silicon film, and the "width axis" may also be neglected due to the electric field's orientation from drain to source. In fact, unlike the illustration in Figure 3.5, in FD-SOI since charge transport mainly happens along the transistor's "length" axis, a 1D simulation can capture both the total charge *and* the average carrier density, provided the virtual "width" dimension is representative of the ion track. Last but not least, as described generically in Section 2.4.2 (page 70), coupling with an external circuit is enabled thanks to the interactive mode of Mentor Graphics' Eldo circuit solver [8], which makes it possible to run a transient simulation step by step with dynamic access to the circuit parameters.

In the following, we first go through the equations that need to be solved (including how to discretize them properly), and shed some light on practical implementation. We then validate CaTHODE against detailed TCAD simulations, and integrate the model inside our Monte Carlo radiation simulator TIARA.

3.3.2 Solver formulation

Physical models

Guided by Occam's razor, we wish to only include in our solver those physical models that are necessary to account for single-event mechanisms in FD-SOI. We thus consider the drift–diffusion transport equations (e.g. instead of more sophisticated hydrodynamic transport):

$$\frac{1}{q}\nabla \cdot \varepsilon \mathcal{F} = p - n + N_{\mathrm{d}} - N_{\mathrm{a}} \tag{3.9}$$

$$\frac{1}{q}\nabla \cdot \mathcal{J}_n = \frac{\partial n}{\partial t} + R - G \tag{3.10}$$

$$-\frac{1}{q}\nabla \cdot \mathcal{J}_p = \frac{\partial p}{\partial t} + R - G \tag{3.11}$$

$$\mathcal{J}_n = q\mu_n n\mathcal{F}_n + qD_n\nabla n \tag{3.12}$$

$$\mathcal{J}_p = q\mu_p p\mathcal{F}_p - qD_p\nabla p \tag{3.13}$$

(3.9) is Poisson's equation, where ε is the dielectric permittivity of silicon, $\mathcal{F} = -\nabla V$ is the electric field obtained from electrostatic potential V, q is the elementary charge, n (resp. p) is the electron (resp. hole) density, and N_{d} (resp. N_{a}) is the donor (resp. acceptor) density assumed to be totally ionized. (3.10) and (3.11) are the electron and hole continuity equations, where R and G are the recombination and generation terms, and $\mathcal{J}_n$ (resp. $\mathcal{J}_p$) is the electron (resp. hole) current density. Finally (3.12) and (3.13) are the drift–diffusion equations for electrons and holes, where μ_n, D_n, μ_p, D_p are their respective mobilities and diffusivities. The driving fields $\mathcal{F}_{n,p} = -\nabla V_{n,p} = -\nabla(V \pm \frac{1}{2}V_{\mathrm{bgn}})$ account for the spatial variations of both the electrostatic potential and the energy gap $E_{\mathrm{g,eff}} = E_{\mathrm{g}} - qV_{\mathrm{bgn}}$, the latter arising because of BandGap Narrowing (BGN) at high doping densities. We account for BGN with the Del Alamo model [122], and the $\frac{1}{2}$ factors imply equal splitting between the conduction and valence bands. We assume Maxwell–Boltzmann statistics for the carrier concentrations (as opposed to the Fermi–Dirac distribution which is more tedious to implement), because our TCAD simulations only showed small differences in terms of radiation currents with or without Fermi statistics. Therefore, the mobilities and diffusivities are related by Einstein's relation: $D = \mu k_{\mathrm{B}}T/q$. On the other hand, note that overlooking bandgap narrowing leads to completely unrealistic predictions for the bipolar amplification mechanism: as was justified in Section 3.1.1 (page 80), with highly-doped implants if BGN is neglected, the minority concentration is off by an exponential factor, yielding unrealistic projections for the parasitic bipolar gain. For orders of magnitude, when deactivating Fermi statistics in TCAD, the collected charge was found to change by no more than 10%, while without BGN it was overestimated by up to 5×. We also model temperature-dependent bandgap and conduction/valence band effective densities of states [133–135]. Finally, note that we define the electrostatic potential with respect to the intrinsic Fermi level, and therefore at equilibrium we have:

$$n_{\mathrm{eq}} = n_{\mathrm{int,eff}} \cdot \exp(qV/k_{\mathrm{B}}T) \tag{3.14}$$

$$p_{\mathrm{eq}} = n_{\mathrm{int,eff}} \cdot \exp(-qV/k_{\mathrm{B}}T) \tag{3.15}$$

Now when outside equilibrium, similarly to the TCAD setup described in Section 3.1.1 on page 80, more specific mechanisms that must be included are:

- *Mobility:* We model the temperature and doping dependence of μ_n and μ_p via the Arora model [136]. We include velocity saturation (i.e. mobility degradation with the electric field) via the Canali model [137], and take Carrier-Carrier Scattering (CCS) into account via the Conwell–Weisskopf model [138,139]. As justified in Section 3.1.1, CCS needs to be included in our solver for correct behavior in the high-LET regime.

- *Recombination:* We consider Shockley–Read–Hall (SRH) recombination depending on the doping level, and Auger recombination depending on the temperature. Auger recombination becomes especially important at high LET.

- *Generation:* We use the Van Overstraeten model [126] for impact ionization. As justified in Section 3.1.1, avalanche effects can increase the overall charge amplification factor. Last but not least, the radiation-induced generation of carriers is accounted for with:
$$G_{\mathrm{rad}}(x,t) = N_{\mathrm{rad}}(x) g_{\mathrm{rad}}(t) \tag{3.16}$$
where $N_{\mathrm{rad}}(x)$ is the space-dependent deposited density (per unit volume) and $g_{\mathrm{rad}}(t)$ is the temporal rate (per unit time) at which deposition occurs.

Finally, we assume electrothermal equilibrium at the contacts at all times. Using l and r subscripts for the left and right sides of the simulation domain, this leads to the following Dirichlet boundary conditions on the primary variables V, n, and p:

$$V(x_l,t) = V_{\mathrm{eq}}(x_l) + V_{\mathrm{a},l}(t) \tag{3.17}$$

$$n(x_l,t) = n_{\mathrm{eq}}(x_l) \tag{3.18}$$

where V_{a} stands for the voltage applied by the circuit. For the sake of brevity, equations in r on the right side and in p for holes were not written.

Spatial and temporal discretization

In order to yield a system of equations that can actually be solved on a computer, the set of Partial Differential Equations (PDEs) above needs to be discretized on a computational grid, both in space and time. The mathematical details of the discretization procedures are given in Appendix C, page 153, here we just wish to mention the main steps taken without resorting to formulas:

- The box discretization method [140] is used to discretize (3.9), (3.10) and (3.11), based on the general non-uniform 1D mesh of variable cross section illustrated in Figure 3.6. Doing so, we can account for the detailed morphology of 28 nm FD-SOI transistors, whose silicon film thickness is 7 nm under the gate, but varies near the spacers and goes to 12 nm in the epitaxy-raised source/drain implants. The electric field $\mathcal{F}$ is easily related to the potential V via a finite-difference formula, but the current densities $\mathcal{J}_n$ and $\mathcal{J}_p$ cannot be obtained from a naive discretization of (3.12) and (3.13): these equations need to be discretized using the Scharfetter–Gummel scheme [141], or else the mesh has to be extremely dense to ensure convergence.

- A backward-Euler approach is used for temporal discretization, for two simple reasons: it is extremely stable (while not the most accurate), and we are anticipating on coupling with an external circuit solver: Eldo's interactive mode allows us to run a transient simulation for arbitrarily small time steps, but it is impossible to "rewind"

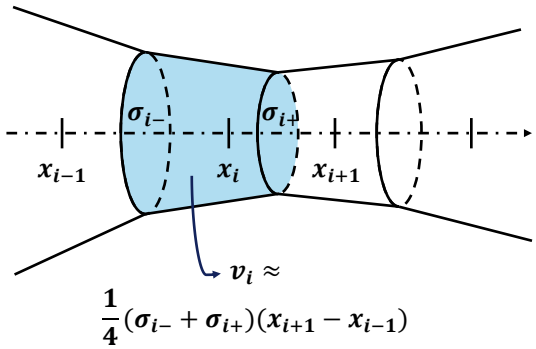

Figure 3.6: Notations for 1D mesh of variable thickness along x.

if we want to subdivide the step further. This prohibits the implementation of a temporal discretization scheme such as the trapezoidal rule, which makes use of "two half time steps" for better accuracy.

- Assembling all equations yields a non-linear system for the potentials and carrier densities, from one time step to the other. The system is solved with the Newton–Raphson method, which ultimately amounts to solving a linear system involving the Jacobian matrix of the equation residuals. Lower–Upper (LU) factorization of the Jacobian is the usual procedure for 1D grids, which leads to two triangular systems solved by forward/backward sweep.

- All of the above allows to propagate in time the space-dependent solution, with extrinsic boundary conditions—but the circuit response depends itself on the device evolution. Since our circuit solver is a black box, we cannot just append the circuit equations to the device equations, and solve the entire system as one single self-consistent entity (which is what mixed-mode TCAD often does). Therefore, additional care must be taken with the device-to-circuit communication: a basic implementation of circuit coupling is to substitute the device with a pure current source whose magnitude is calculated by the device solver, then let the circuit solver refresh the terminal voltages, and then compute a new value for the device current, and so forth. Such a coupling method was tested but exhibited mediocre stability, typically requiring time steps below 0.1 ps at all time for convergence. Based on [142], a much better convergence radius was then obtained by replacing the device with an equivalent current source in parallel with an equivalent conductance, as depicted in Figure 3.7 (for a benchmarking circuit discussed after). Just like the current value, the conductance value is refreshed at each time step by the device solver (with negligible calculation overhead), then handed over to Eldo for circuit simulation. This allows to raise the time step to values close to the picosecond range during the impact—and much higher after—with very limited accuracy loss.

3.3.3 Implementation details

Our model is written in about 5 000 lines inside TIARA's C++ code. During the transient run, communication with the circuit solver is implemented by redirecting the output flux of our program to the input flux of Eldo in interactive mode, as presented generically in Chapter 2 (Section 2.4.2, page 70): using "pipe" system calls, at each time step we

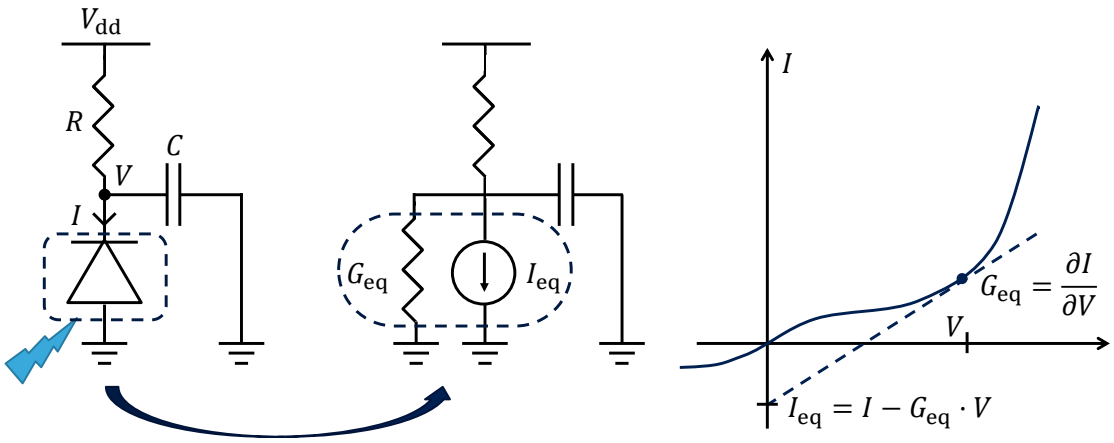

Figure 3.7: Circuit coupling with the simulated device's equivalent conductance and current source inside the electrical solver.

are able to communicate the equivalent current and conductance (device-to-circuit link), and retrieve the updated terminal potentials (circuit-to-device link). Such an implementation of device–circuit connection was already presented in [103]. Although some of the syntax is necessarily Eldo-specific, this solution can be ported to other electrical solvers as soon as they offer a command-line execution mode allowing arbitrary time increments and dynamic access to the circuit parameters. Alternatively, other works have demonstrated the use of Verilog-A models to perform a similar task [68]. Note also that before the transient simulation can be run, the solver first needs to find the equilibrium situation, or zero-bias steady state, by solving (3.9) self-consistently with (3.14) and (3.15), page 89, starting from an initial guess based on charge neutrality. It must then perform a quasi-stationary ramp-up (taking out all $\partial f/\partial t$ terms) to reach a steady state under finite bias, so that the device is ready to be irradiated.

Now, several optimizations are needed to ensure fast execution:

- First of all, the linear system obtained for the Newton method needs to be built before it can be inverted. Calculating all of the system's protagonists is very lengthy if nothing is done to speed-up the expression evaluations. To tackle this issue, we adopt a pragmatic, hybrid strategy by caching the "clearly re-usable" arrays. For instance, the SRH recombination term is calculated with a fraction N/D on each grid point, so storing the N and D arrays allows for faster evaluation of $(N'D - ND')/D^2$ in the Jacobian calculations later on. This requires some hard-coding of the dependence tree in the expression evaluations, to avoid refreshing values in the wrong order. Overall, we observed a near 10× speed-up when going from a careless evaluation methodology to this hybrid solution. Therefore a complete handling of expression trees—which would be required for optimal lazy re-evaluation—was deemed unnecessary.

- Second of all, the system needs to be inverted. For this we make use of Eigen [15], a very complete library for linear algebra. It allows to analyze the sparsity pattern of the Jacobian matrix to accelerate its LU factorization afterwards: the non-zero entries do vary but their indices do not, so the sparsity analysis can be performed once and for all beforehand. We observed a 40% performance improvement with this feature—which ironically made the system faster to invert than to build. Furthermore, Eigen enables many compiler optimizations to vectorize the calculations, yielding an

additional 3× speed-up. We also implement some adaptive time step control tailored to our radiation context, with time step relaxation after the impact for instance. The end result of all these optimizations is that an impact takes about 1–3 seconds to run, depending on its toughness: this perfectly meets our speed requirements for SER simulation.

3.3.4 Model validation

Comparison with 3D TCAD

As a benchmark to validate our model in FD-SOI technology, we compare the results of our 1D solver with more detailed TCAD simulations. The latter are performed with a similar setup as presented in Section 3.1, by striking a 3D nMOS structure connected in mixed mode to a pMOS and a load capacitance, to simulate the context of an inverter in a CMOS network (Figure 3.8a). Our 1D solver simply takes as input the 1D doping profile of the struck nMOS along its channel axis. A very important note is that in our model, we do not "erase" the nMOS in the circuit simulation, unlike what was shown in Figure 3.7 for a simplified circuit: we cannot realistically hope for our solver to reproduce the full behavior of a MOSFET since the influence of the gate cannot be natively treated in 1D. Therefore, for the equivalent current and conductance calculations we subtract the steady-state leakage predicted by our solver, and keep the MOSFET in parallel with those circuit elements (see Figure 3.8b). In other words, we leave the silicon-calibrated SPICE transistor in charge of the "normal-operation" current, and our solver is only responsible for the radiation-induced current. From a device-level perspective, this means that our simulator is accountable for the motion of excess carriers but not the "background" electrons and holes.

Figure 3.9 displays the current and voltage waveforms obtained for impacts at the drain-channel junction of the transistor (where the electric field is highest), for various LETs. For comparison we also include results from BA-Ba, the behavioral model developed previously, where the current waveform was simply a double-exponential pulse, of integral charge given by an LET-dependent bipolar amplification modulated by a

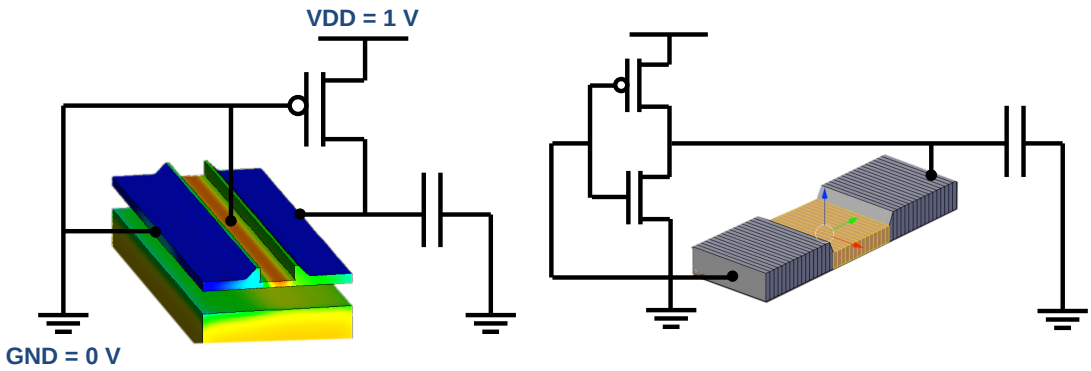

(a) (b)

Figure 3.8: (a) Mixed-mode 3D TCAD simulation of an inverter (nMOS strike) – (b) Simulation setup with our single-event current model for the same radiation impact; the SPICE nMOS is kept in the netlist and the DC leakage predicted by our 1D solver is subtracted to avoid inconsistencies.

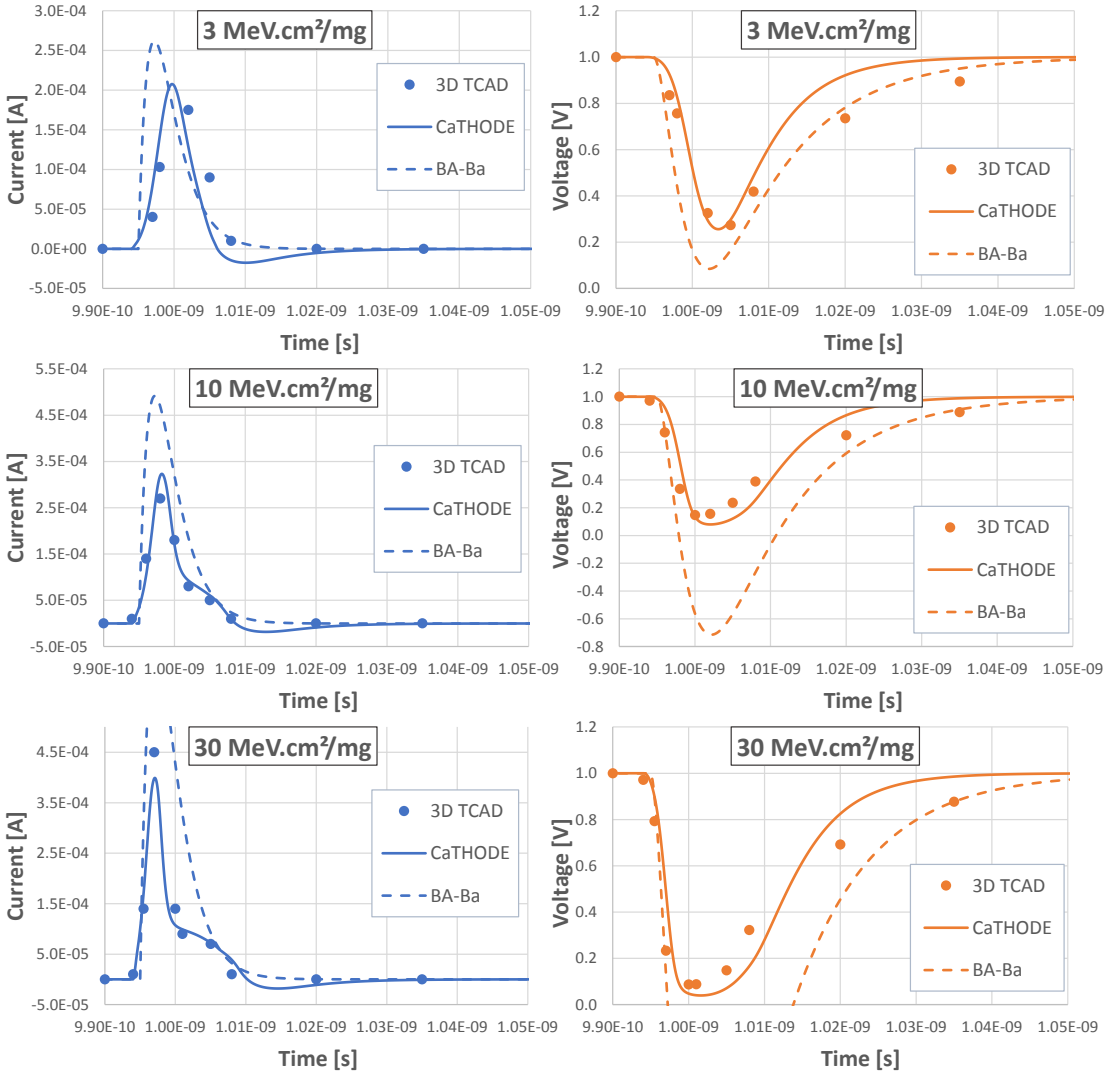

Figure 3.9: Comparison between our FD-SOI response models and mixed-mode 3D TCAD simulations (impacts on nMOS inside an inverter across LETs).

position-dependent collection efficiency. As can be seen, the time-dependent effects are well accounted for by our new model, which natively enlarges the current pulse in high injection. On the other hand with the simple model, at high LET the decoupled current source is way off since it does not vanish when the potential drops, thus causing unphysical values for the voltage. This is not an issue for SEU simulation which is merely interested in whether the voltage does drop or not. But SET simulations would be quite coarse with BA-Ba, since the voltage pulse duration is mispredicted.

Mention should be made that the model developed here is not exempt from calibration: the TCAD simulations shown here use hydrodynamic transport equations with more complex mobility modeling (especially thin-layer mobility with interface degradation), beside the models already listed in Section 3.3.2, page 89. Thus some adjustments are necessary, mainly by playing on the carrier mobilities and lifetimes. Note however that obtaining a good match with TCAD in not an end *per se*: silicon measurements—e.g. comparison with SEU cross sections on sequential cells as shown later on—still have to remain the main judge for Model-Hardware Correlation (MHC).

Discussion on circuit coupling

As a mere sanity check, in Figure 3.10 we plotted the waveforms obtained with our model for a 10 MeV·cm^2/mg impact on a p-n$^+$ structure with mesh dimensions representative of a bulk-Si junction as depicted in Figure 3.7 (page 92), with $V_{\mathrm{dd}} = 1.2$ V, $R = 10$ kΩ, and $C = 1$ fF. As can be seen, our 1D solver perfectly captures the so-called "plateau" effect [112, 113], which arises in modern bulk technologies when the time constants for charge transport and circuit response start overlapping. The n$^+$ voltage is clamped at $-V_{\mathrm{bi}} = -0.8$ V, while the current lingers at the precise value of $(V_{\mathrm{dd}} + V_{\mathrm{bi}})/R = (1.2 + 0.8)/10\,000 = 0.2$ mA: if the current was to exceed this value, the n$^+$ potential would drop even further and the junction would become truly forward biased, thus it would start driving a massive current in the opposite direction that would actually draw the potential back up very quickly. As a matter of fact, this phenomenon can be observed when intentionally using overly relaxed time steps during the plateau: the forward-bias current restores a dramatically high potential on the n$^+$ node, which causes spurious oscillations that sometimes stabilize, or sometimes can cause the simulation to diverge.

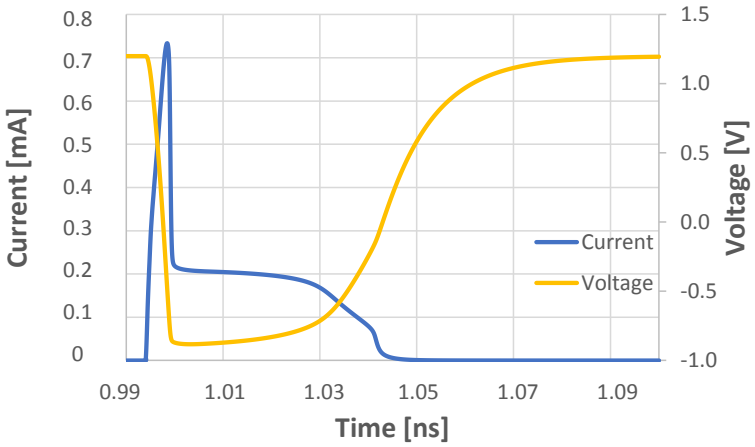

Figure 3.10: Sanity check for circuit coupling: plateau effect for a 10 MeV·cm^2/mg impact on a mesh representative of a bulk-Si p-n$^+$ junction connected like in Figure 3.7, page 92. The plateau lasts for as long as there are charges to collect that have not recombined, and then the voltage recovers with an RC time constant.

On the other hand as can be seen in Figure 3.9, in FD-SOI the plateau effect due to circuit coupling does arise, but in a less pronounced way than in bulk Si technology. This was expected given the reduced active silicon volumes, providing a much smaller reservoir for continued charge collection over time. More important than the *current* plateau, however, is the fact that the V_{ds} *voltage* of the MOSFET is clamped at 0, much like it was clamped at V_{bi} for a junction alone. This is of great relevance to properly capture the radiation behavior of devices connected in series. This will be further elaborated upon in Section 4.2.1, page 110.

3.3.5 Integration in TIARA

In order to enable smooth integration of the new collection model within TIARA, several software upgrades were needed on both sides. Beside general flow unification tasks to

connect the APIs together, we had to dynamically generate the inputs for the new CaTHODE `Collector`:

- First, automatic generation of the 1D meshes had to be implemented, based on the impacted transistors. Geometrically, this is done by splitting the geometry that the `Builder3D` module originally delivers, with a certain targeted grid step typically around 1 nm. In the case of two transistors sharing a common diffusion area, we generate two simulation domains with a shared boundary condition, by simply splitting the implant geometry halfway. Note that in such aggressively integrated technologies as 28 nm FD-SOI, Design Rule Check (DRC) checks forbid having more than two FETs sharing an implant. This avoids complications as to domain partitioning—moreover for star-like connections, the 1D transport assumption would be very challenged.
- After the 1D meshes have been created for the impacted devices, they need to be assigned doping profiles as well; since all possible gate lengths can exist in the simulated circuit, this is better done with fitting functions rather than table look-up. Based on the TCAD deck used previously, active dopant concentrations were extracted along the channel axis for nMOS and pMOS transistors from $L_{\mathrm{min}} = 30$ nm to $L_{\mathrm{max}} = 10$ µm, in all of the technology's V_{t} flavors. The obtained profiles were then fitted with sigmoid curves describing the slopes of the donor/acceptor densities around the junctions. For each device type, this results in a 13-parameter set encoding the doping profile as a continuous function of gate length (Figure 3.11a).
- Finally, for the transient transport simulation to begin, the outputs from TIARA's `Raytracer` module need to be converted into a carrier generation term. This is done by splitting the particle track segments along the cells of the 1D mesh, and in order to avoid singularities, this raw charge deposition is smoothed out with a Gaussian convolution kernel (Figure 3.11b). To some degree, this allows to emulate the ion track structure, as will be discussed in Appendix B. Now as a mere illustration of the transient evolution simulated by CaTHODE, in Figure 3.11c we can see the collapse of the electric field inside an SRAM's pMOS being struck, resulting in an SEU.

3.4 Chapter conclusion

This chapter has been devoted to transistor-level simulation and modeling of single-event effects in 28 nm FD-SOI. First, the technology's response has been investigated by means of detailed 3D TCAD simulations on elementary MOSFETs. The main outcome of our TCAD study is that both the sensitive volume for charge collection, and the parasitic bipolar amplification, are small in UTBB SOI. The most sensitive area lies at the drain–channel junction of the device, and thanks to the thin silicon film, this leads to small collected charges in the femtocoulomb range, even at high LET. The parasitic bipolar gain was shown to be below 3 under nominal bias conditions and for SEE-relevant LETs. This means that at circuit level, the technology is expected to be very resilient to soft errors.

These two main specificities of 28 nm FD-SOI were then encoded in a behavioral SEE model called Bipolar Amplification Basics (BA-Ba) for use in our Monte Carlo SER simulator. BA-Ba computes charge collection by assuming an LET-dependent bipolar amplification, and a space-varying collection efficiency. After computing the charges collected by sensitive impacted MOSFETs, BA-Ba injects them as double-exponential

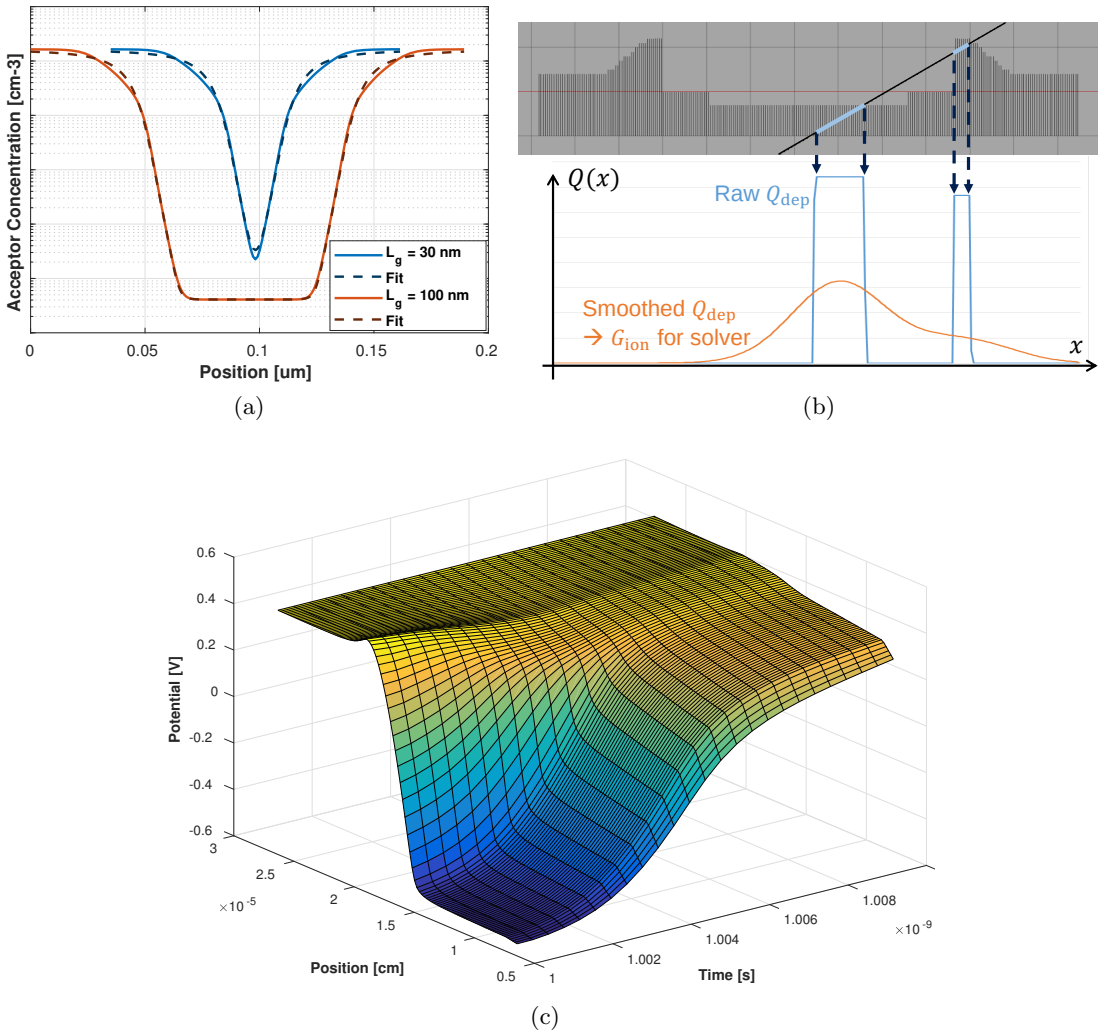

Figure 3.11: CaTHODE collection model integration in TIARA: (a) Doping profile analytical reconstruction w.r.t. pMOS gate length – (b) Carrier generation term calculation from the ray-traced ion track – (c) Electrostatic potential inside pMOS during an SEU-inducing impact.

current waveforms of corresponding integral charge. The model parameters are based on values tabulated from TCAD, and further ahead, calibrated on heavy-ion SRAM SEU cross sections. Furthermore, it was shown that since BA-Ba does not implement the circuit's feedback on the charge collection process, it shall mainly be reserved to SEU investigations.

Finally, we have developed a time-dependent single-event model for FD-SOI named Carrier Transport with Humble One-Dimensional Equations (CaTHODE), which is interfaced with an electrical solver. The model is based on direct resolution of the semiconductor equations on a one-dimensional grid, making it computationally efficient. Furthermore, it was shown that the model is very accurate when applied to 28 nm FD-SOI devices: owing to the ultra-thin silicon film in that technology, the "1D" modeling choice is justified, and the model was shown to compare very well with mixed-mode 3D TCAD simulations, properly capturing the dynamic effects of circuit feedback on collection. CaTHODE has been integrated within our Monte Carlo radiation simulation platform TIARA, by automatically generating 1D grids, doping profiles, and carrier generation terms for the impacted devices. It should also be mentioned that during this PhD,

CaTHODE was successfully applied at the 22 nm node in FD-SOI. With some adaptation, the model could be applied to FinFET technology as well: using simplified 3D diffusion in the bulk beneath the fin, and looking at carriers collected at its base (i.e. at the STI openings), an effective generation term inside the fin could be derived. Then, this generation term would be the entry point to a 1D transport simulation as shown in this work. Note that this implicitly assumes that collection inside the fin does not strongly affect transport in the bulk beneath. This seems to be a reasonable assumption, since the electric field is mostly confined within the channel.

From a physicist's perspective about CaTHODE, compact analytical models might often be preferred for the immediate insight they can offer, while brute-force numerical approaches are sometimes regarded as merely "number-crunching" routines making it hard to build one's physical intuition. However in this case we find that, owing to the solver's great speed, having the ability to get an instantaneous answer to any thought experiment—in 1D—is also extremely valuable and helps us gain a deeper understanding of the underlying mechanisms.

Chapter 4

Logic-cell level simulation of radiation effects

This fourth chapter is dedicated to the study of single-event effects at circuit level. Using the innovative response models that were developed in Chapter 3 for FD-SOI technology, and previously existing models in bulk, we perform TIARA Monte Carlo simulation studies on several types of cells. This allows us to discuss the various contributors to the soft-error rate in digital CMOS logic gates. Our presentation also draws from a host of experimental data acquired in beam test campaigns; the simulation analyses are not only complemented by, but actually derived from test results, since our SEE models always need to go through a Model-Hardware Correlation (MHC) phase to calibrate the simulations against silicon data. The measurements presented here were not produced by the author of this thesis, but such test results are thus absolutely pivotal to our understanding of the physical mechanisms, both as theorists or experimentalists.

In Section 4.1 we first discuss the response of FD-SOI SRAM bit-cells and sequential logic cells such as flip-flops. Using our SEE models tailored to FD-SOI and inserted in TIARA, we establish the key features of 28 nm UTBB SOI in terms of SEU robustness. We then present simulations on combinational gates carried out both in 65 nm bulk silicon and 28 nm SOI technologies, namely the two main ST space platforms of interest in this work. Indeed, single-event transients affecting combinational cells have become a major concern in advanced technologies; since SEUs can be addressed rather effectively by standard hardening mechanisms, SETs constitute the main soft-error hazard remaining in typical SoCs—at least as far as digital CMOS blocks are concerned. How they contribute to the overall SER, and how they can be mitigated, is the topic of Section 4.2. Finally, in Section 4.3 we examine the behavior of both planar bulk and FD-SOI technology architectures in various space environments. By folding our energy-dependent characterizations against simulated particle spectra, we discuss the on-orbit upset rate of present-day technologies along with concrete implications for real-life applications.

4.1 Single-event upsets in FD-SOI technology

Let us first begin with the SEU response of FD-SOI technology, by studying irradiation and simulation results of SRAMs and DFFs. Some of the content of this section is adapted from [13]. Note that our presentation is centered on the 28 nm process node,

but our conclusions would hold for other advanced UTBB technologies. In particular, 22 nm measurements and simulations were produced within the timeline of this PhD, the overall outcome being that both technologies display quite similar behavior radiation-wise.

4.1.1 Memory cells in 28 nm: experimental results and calibration of TIARA collection models

Heavy ion irradiation results on SRAM cells

In order to qualify the digital building blocks of FD-SOI, several test vehicles were designed and manufactured in STMicroelectronics' 28 nm UTBB SOI. GDS and micrograph views for one of them, called SERVAL, are given in Figure 4.1: the test chip embeds several flip-flop registers and SRAM cuts, as well as microprocessors with various implementations (reference, low voltage, radiation hardened). Heavy-ion testings were performed at RADEF, Finland [46, 47], in compliance with ESA/ESCC basic specification No 25100 [50]. A subset of the 9.3 MeV/amu cocktail was used, namely nitrogen, neon, iron, krypton and xenon ions giving rise to LETs of 1.83, 10.2, 18.5, 32.2 and 60.0 MeV$\cdot$cm^2/mg. Three parts were irradiated with fluences of about 10^7 ions/cm^2 on each run and, except for the SEL runs, nominal stress conditions were used, i.e. at room temperature and with the memory arrays' supply voltage at $V_{\mathrm{dd,mem}} = 1.0$ V.

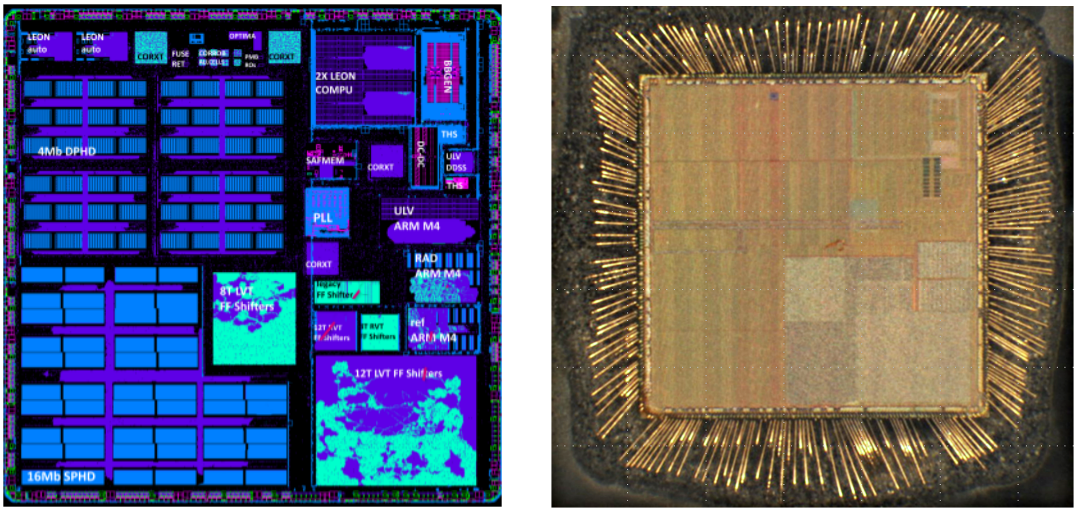

(a) (b)

Figure 4.1: SERVAL test vehicle irradiated for 28 nm FD-SOI qualification: chip layout (a) and micrograph after package etching (b). The SRAM cuts occupy the left half of the design.

The overall outcome of this test campaign is a very low sensitivity exhibited by the FD-SOI technology under heavy ions. Previous publications had already reported the outstanding SEU hardness of 28 nm FD-SOI under terrestrial particles such as alphas and neutrons [71,75]; these irradiation results confirm the technology's hardness in space environments as well, with heavy-ion cross sections routinely achieving 2 decades of reduction when compared to bulk Si counterparts. As was investigated with TCAD simulations in Section 3.1, this extreme robustness of FD-SOI is primarily due to two factors: first of all the very small volume for ion-induced charge *deposition*, thanks to the 7 nm-thin active silicon film being enclosed by the BOX. Second of all, parasitic

bipolar amplification is quite limited in this technology, leading to *collected* charges as small as 0.1 to only a few fC.

In the rest of this section we will focus on single port bit-cells whose schematic and layout is similar to those depicted in Figure 1.15a on page 24. The experimental heavy-ion cross sections collected in the beam tests for one such cell are shown in Figure 4.2 along with a four-parameter Weibull fit. As an element of comparison, the cross section for a 28 nm "bulk-equivalent" bit-cell was also plotted (identical area and similar design, give or take some DRC alterations). As can be seen, the FD-SOI cell is about two decades less sensitive than its bulk counterpart: the bulk bit-cell exhibits an asymptotic cross-section of about twenty times the cell area, because charge carriers are free to diffuse on very long ranges. Note that we are reporting cross sections from bit flip counts, not "event" counts, therefore accounting for MCU multiplicity. Error maps for the MCU patterns are not shown here but do confirm that on average, each impact triggers a twenty-cell upset in the bulk bit-cell array at 85 MeV$\cdot$cm^2/mg. On the other hand, the FD-SOI limiting cross section σ_{sat} is only about 15% of the cell area, thanks to the full dielectric isolation provided by the BOX. In bulk technology, the most sensitive areas are the drains of off transistors (with a preponderant vertical field), and for an LET as low as 1.5 MeV$\cdot$cm^2/mg the cross section already equals this cumulated surface. Then as the LET increases, the cross section expands much beyond this "primary" sensitive surface, as mentioned above. However in FD-SOI the situation is very different: the sensitive areas are rather located by the off-transistor channels featuring a strong lateral field—and although a few upsets are measured at the lowest LETs, it takes an LET of about 15 MeV$\cdot$cm^2/mg for the cross section to be equal to this cumulated surface. Then at higher LET, the cross section increases by no more than 50% unlike in bulk. These experimental area measurements, albeit somewhat indirect, corroborate the TCAD study of sensitive volumes that was carried out in Section 3.1.

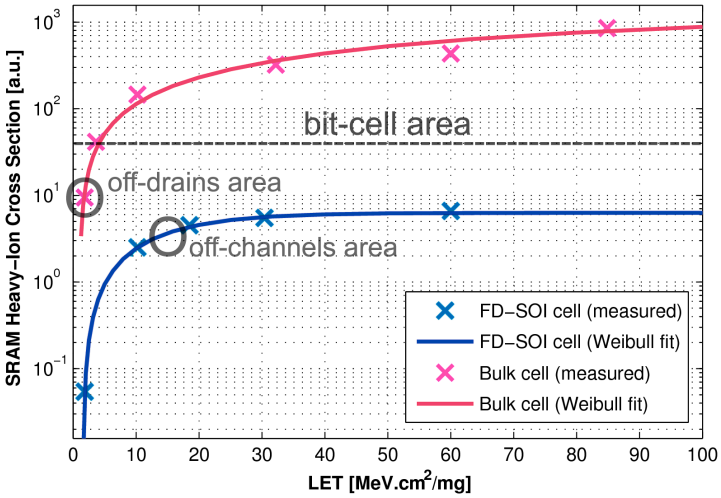

Figure 4.2: Experimental heavy-ion cross sections of a single-port bit-cell, exhibiting FD-SOI's intrinsic robustness compared to bulk Si.

Calibration of the behavioral collection model (BA-Ba)

With this in mind, we can proceed to calibrating the FD-SOI collection models used by TIARA. We set up our Monte Carlo simulations to mimic the experimental conditions,

with heavy ions of various LETs vertically incident on bit-cells at V_{nom}. We begin with BA-Ba, our behavioral collector with precomputed bipolar amplification. The constitutive equations (3.7) and (3.8) we proposed (page 85) contain parameters whose orders of magnitude and relative trends are suggested by our TCAD simulations, but can remain somewhat free for fine tuning against experimental data:

- Precise numbers for $\beta(LET)$—the LET-dependent bipolar amplification of nMOS and pMOS—are obtained to match the observed threshold LET while keeping a value of one at high LET. Concretely, we use a linear transformation on the raw values obtained in TCAD to meet those two criteria.
- Defining the spatial dependence of $\eta(x)$, or the non-uniform collection efficiency along the 'L_{g}' axis (gate length), truly requires nanometric resolution around the drain–channel junction in order to match the onset of experimental heavy-ion cross section curves. Then, the point in space where this collection yield goes to zero dictates the maximum sensitive area seen under normal incidence, hence it is taken to reproduce the measured saturation cross sections.

The other parameters in BA-Ba, namely the rise and fall times τ_{r} and τ_{f} of the double-exponential current pulse, are kept as discussed in Section 3.2. Our calibration procedure is illustrated in Figure 4.3, with the $\beta(LET)$ and $\eta(x)$ functions being fine-tuned to match experimental data on two bit-cells, a 0.120 $\mu\mathrm{m}^2$ Single-Port High Density (SPHD) cell and a 0.197 $\mu\mathrm{m}^2$ Single-Port Level-One Cache (SPL1$) cell. Note that beam test results are presented for other test chips than SERVAL mentioned previously, namely QLIB and QSRAM dedicated to memory cells qualification and irradiated at Université Catholique de Louvain (UCL), Belgium [143] and RADEF. Although there is some dispersion in the results due to different silicon, it can be seen that firstly, SPHD exhibits a lower threshold LET because of its higher density, enabling low-LET calibration—to adjust especially β against LET_{th}. Secondly, SPL1$ has higher saturation cross section because of its larger area, providing us with high-LET information to tweak η against σ_{sat}.

Calibration of the advanced response model (CaTHODE)

Next, we proceed with CaTHODE, or our response model based on a 1D drift–diffusion solver with dynamic circuit coupling. Unlike the more empirical model BA-Ba, CaTHODE's tuning knobs are physical parameters inside the semiconductors equations. As discussed in Section 3.3.2 on page 89, the main levers we can use are first of all, bandgap narrowing, whose magnitude can be controlled by V_{bgn} in (3.2) (page 81) with dramatic effect on bipolar amplification. Second of all, the carrier mobilities offer fine-grain adjustment possibilities, through carrier-carrier scattering and velocity saturation models. This is what we show in Figure 4.4, with the impact of the BGN and mobility parameters on heavy-ion cross sections for the SPHD bit-cell.

At this point, we may consider the error maps in Figure 4.5 revealing the sensitive areas around threshold and at saturation. The silicon film thickness t_{Si} we simulate (see Figure 2.3d on page 52) is represented in shades of green; as can be seen, the regions where SEUs first appear result from a competition between t_{Si} (maximum charge deposition within the epitaxy-raised implants) and the electric field (maximum pair separation at the drain–channel junction): the initial sensitive region is not necessarily located right below transistor gates, as is often reported in SOI [11]. At high LET, the

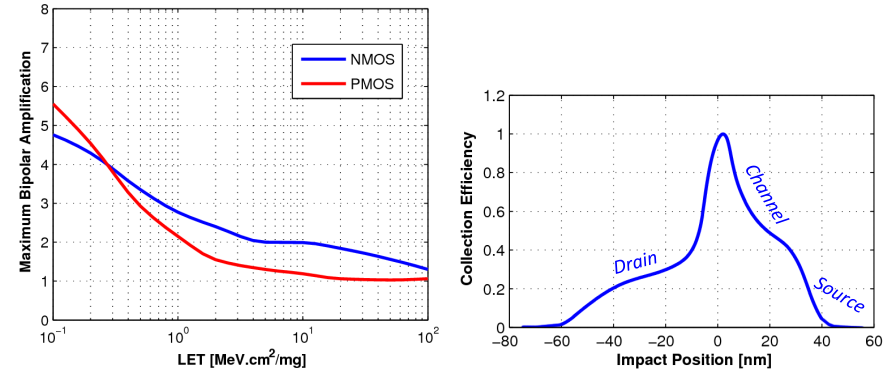

(a) (b)

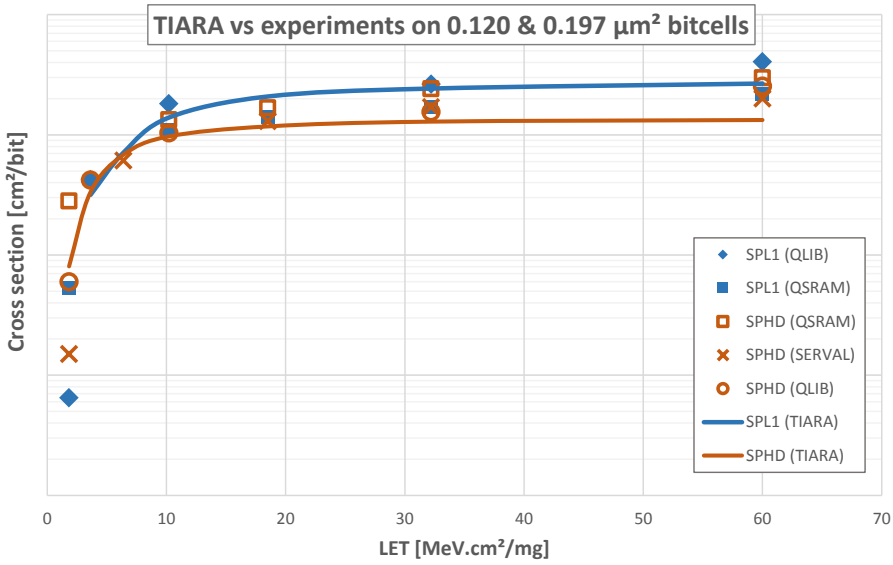

(c)

Figure 4.3: Calibration of TIARA's behavioral FD-SOI collection model BA-Ba from 28 nm SRAM heavy-ion data at normal incidence: (a) LET-dependent bipolar amplification – (b) Space-varying collection efficiency ($L_g = L_{min}$) – (c) Resulting cross sections for two singe-port bit-cells.

sensitive area comprises the channels and drains of the off-state FETs, but does not extend much in the source implants: at these nodes, charges are evacuated harmlessly by the supply and ground rails. Such visual insight is extremely powerful to build one's physical intuition of the mechanisms at work—not only to the radiation expert, but also for knowledge transfer with circuit designers for instance.

Regardless of the collection model, note the importance of having heavy-ion results for our MHC procedure: this allows to unveil the technology's response to pure excitations (i.e. with monoenergetic radiation) spanning a large spectrum. Hence, this is as close as we can get to an exhaustive characterization of our cells—and single-port SRAMs are simple to study because of their unique electrical state (give or take symmetry). The main limitation is that during test campaigns, the DoE usually does not include many angles: test conditions are typically at normal incidence only, or at best, at a few tilt angles e.g. 0°, 45°, 60° (vertical, halfway down, and $2 \times LET_{eff} = LET/\cos\theta$), with 0° and 90° roll (aligned on the MOSFETs' W or L). Therefore, comparison with experimental proton results, as shown in Figure 4.6 with measurements done at PSI [48, 49], is also of great help in our model adjustment phase: although the spectrum of

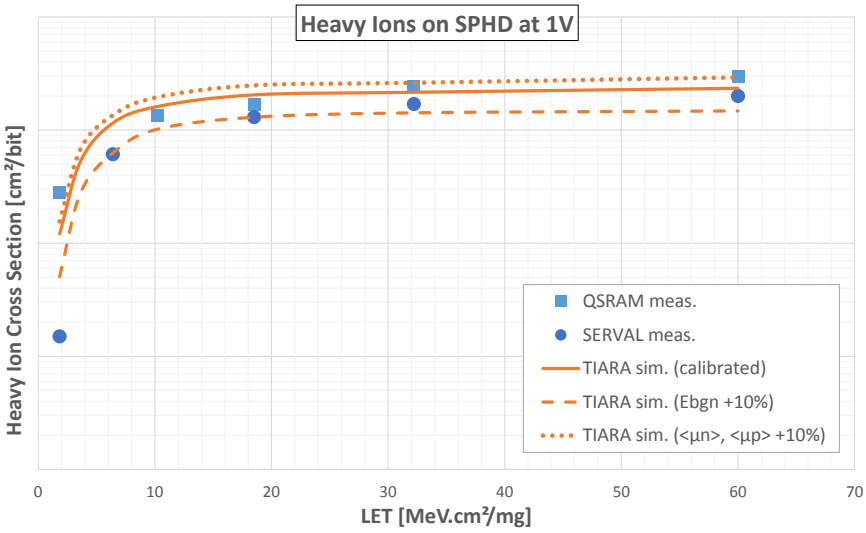

Figure 4.4: Calibration of TIARA's advanced FD-SOI collection model CaTHODE.

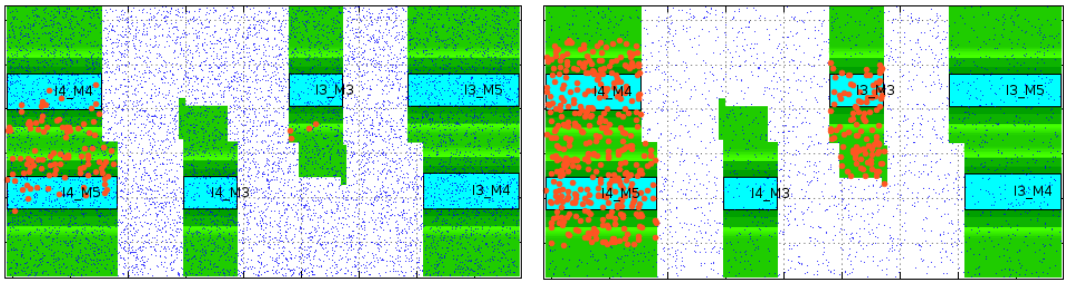

(a) (b)

Figure 4.5: Simulated heavy-ion error maps on SPHD. SEUs in red dots, impact points otherwise in blue. FET channels in cyan, active Si otherwise in shades of green – (a) $7.5{\times}10^{12}$ cm^{-2} nitrogen ions at 1.83 MeV·cm^2/mg – (b) 2.5×10^{12} cm^{-2} xenon ions at 60 MeV·cm^2/mg.

nuclear products is spread out in LET even for monoenergetic protons, this provides us with precious information about the angular response of the technology. Note that the low-energy part of the proton cross section curve is currently not well reproduced in our simulations, with TIARA cross sections being optimistic below 30 MeV. This is because cells with low critical charges such as SRAMs have non-zero sensitivity to low-energy protons. In Section 4.3 we will discuss the implications of Proton Direct Ionization (PDI) on the overall SER in orbit, and how to realistically handle PDI in TIARA.

On a last note in this section on models calibration, let us mention that for our terrestrial applications, FD-SOI collection models were also validated against alpha-particle SER, as reported in [10]; α-rays trigger impacts in nearly the full hemisphere of incoming directions (only grazing particles are excluded due to attenuation in the BEOL), and this helps investigate the angular response of FD-SOI. Furthermore, alphas have low LET, and while this means that they only characterize a small portion of the spectrum, this also leads to the α-SER being strongly voltage dependent in FD-SOI[1]. Thus, the steep SER-versus-V_{dd} curves reported in [10] help us ensure that the electrical response is also

[1] Q_{crit} scales linearly with voltage both in bulk and fully-depleted technologies, as explained on page 24; however in SOI, Q_{coll} is determined by the rigid sensitive volume, whereas in bulk, depletion regions can extend further, as discussed on page 18. Hence the stronger voltage dependence of the SER in FD-SOI.

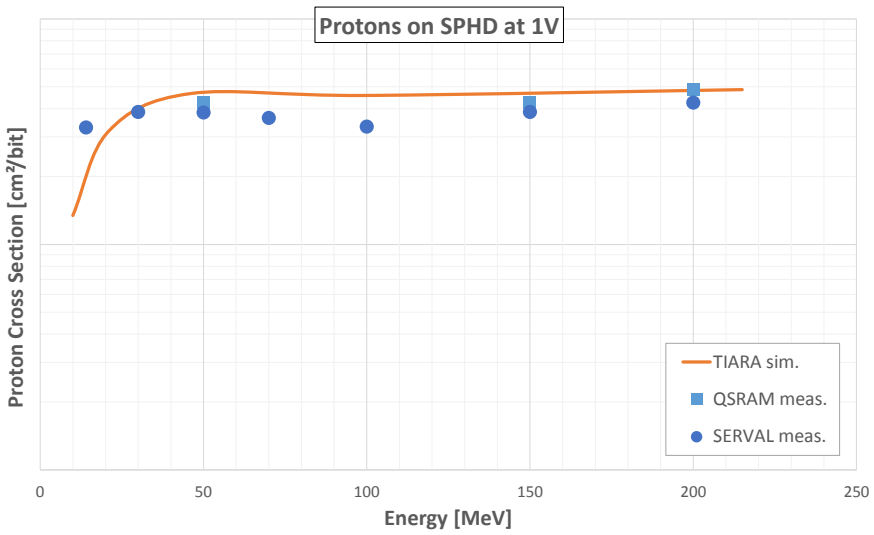

Figure 4.6: Proton cross section results corroborating the simulated angular response.

well captured in our simulations. Overall, the extensive collection of experimental data gathered over many test vehicles, in a variety of operating conditions and with several particle sources, gives us very good confidence in our SEE modeling approach in FD-SOI. As we will discuss in the next sections, both BA-Ba and CaTHODE, when used in their respective validity ranges, allow us to perform predictive TIARA simulation on complex circuits and in various particle environments. This is when Monte Carlo SER simulation capabilities truly become relevant, virtually enabling our DoEs beyond monoenergetic, or unidirectional, or single-operating-point test plans.

4.1.2 Usage of TIARA for 28 nm sequential logic characterization

Throughout the course of this PhD, TIARA was used on several occasions to perform SER estimations on memory cells (with calibrated simulations down to 22 nm and more extrapolative projections in 14 nm); however, on mature technologies, SRAMs are characterized nearly exhaustively, which radiation-wise means that measurement points are readily available for several particle types, operating points, and for most—if not all—compilers in the platform's offer. Thus, for memory cells at well-established process nodes, there are usually few gaps to fill with SER simulators such as TIARA. On the other hand, standard cell libraries often contain dozens to hundreds of logic cells, all of which cannot see particle beams within realistic time-to-market; a handful of them are cherry-picked for representativity, providing silicon data on a sparse subset of the library. This is precisely when we can hope for our simulator to connect the dots, and indeed TIARA was intensively operated on 28 nm FD-SOI sequential logic cells. This includes flip-flops and latches of standard or high-speed architectures, with different design options (reset/enable pins, Design For Testability (DFT) or scan inputs, output drive), and ranging from unhardened to radiation-immune through a variety of robustness levels. For our presentation in this section, we choose to work out just one example with a full-spectrum characterization of an unhardened D-type flip-flop, for which partial experimental measurements were available.

Our object of study is an unhardened master–slave DFF with reset and scan input

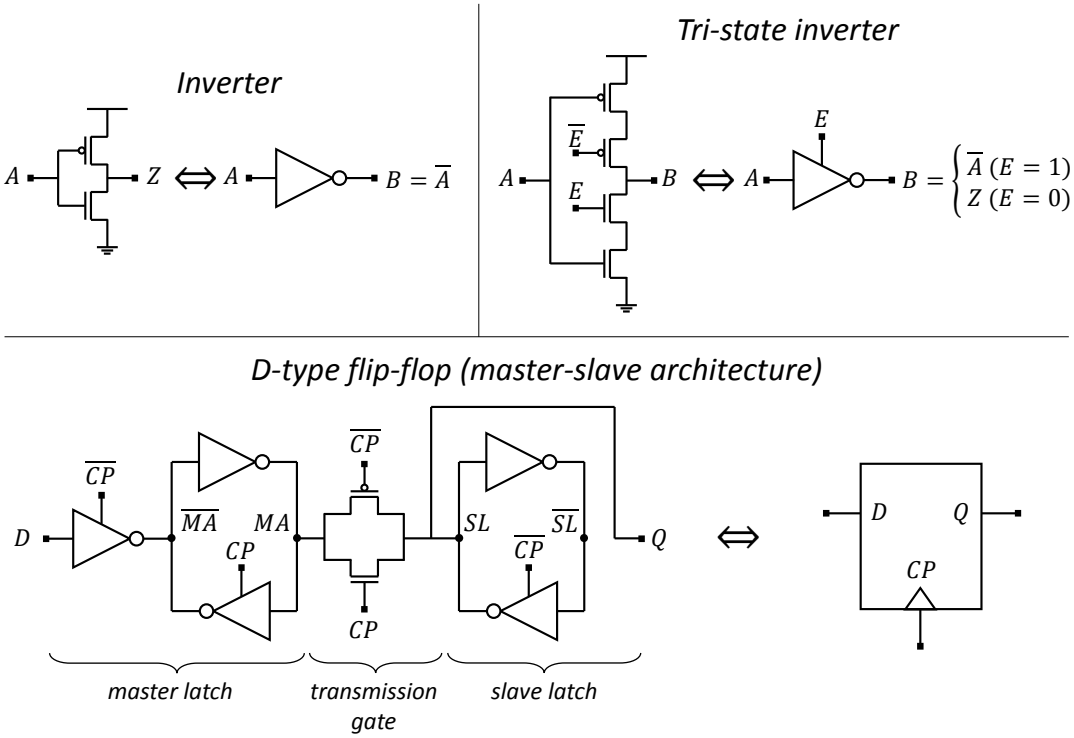

Figure 4.7: Schematic of a master–slave D-type Flip-Flop (DFF).

pins, Regular Voltage Threshold (RVT) flavor[2] and a low output drive current. Its simplified schematic (without the reset and DFT functions) is drawn in Figure 4.7: both the master ($\mathrm{MA}, \overline{\mathrm{MA}}$) and slave ($\mathrm{SL}, \overline{\mathrm{SL}}$) latches are comprised of a forward inverter and a feedback tri-state inverter controlled respectively by the clock pin CP and its complement $\overline{\mathrm{CP}}$. Thus when $\mathrm{CP} = 1$, the slave latch is transparent while the master is memorizing, and *vice versa*. Considering also the transmission gate which is statically conductive, and isolating upon clock transitions, this typical architecture realizes a positive-edge triggered flip-flop, transmitting input D to output Q on the tick of the clock. Statically, there are eight possible states to this DFF depending on the state of the input pins and the internal value stored by the latch that is memorizing: $\{\mathrm{D} = 0, \mathrm{D} = 1\} \times \{\mathrm{CP} = 0, \mathrm{CP} = 1\} \times \{\mathrm{SL/MA} = 0, \mathrm{SL/MA} = 1\}$.

The cell was embedded and experimentally qualified in the SERVAL test chip presented in previous section, under protons and heavy ions respectively at PSI and RADEF. More precisely, the flip-flop is instantiated in a shift register or "shifter", and BIST circuitry is embedded with the shifter: in Figure 4.8, the external FPGA just needs to provide a static pattern, over which the on-chip pattern generator repeatedly loops to insert a dynamic sequence in the shifter; then the on-chip error counter performs all the at-speed comparisons with the reference pattern, alleviating the FPGA of high frequency operations. This allows to perform high-speed runs, typically up to 1 GHz during exposure. The advantage of such a dynamic test protocol is that potential SET mechanisms can be observed, the drawback being that the overall SER or cross-section measurement is a mixture of several electrical states; for instance, even a solid-0 dynamic pattern excites both (D0, CP0, SL0) and (D0, CP1, MA0). During the experiments, a '0011' pattern was used, giving equidistribution among the eight possible static states.

[2]The other V_t option in ST's 28 nm FD-SOI is Low Voltage Threshold (LVT), achieved with lighter doping in the transistor's body.

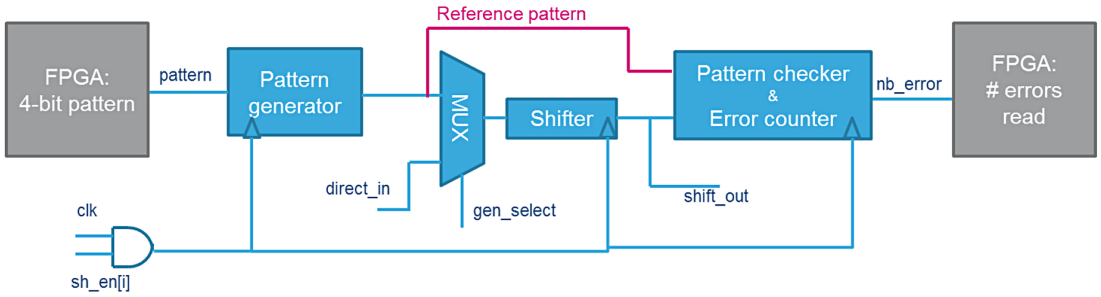

Figure 4.8: Built-In Self Test (BIST) structure for on-chip shifter monitoring (on-chip blocks in blue, off-chip in gray).

Experimental and simulated heavy ion results are given in Figure 4.9. Experimentally, no significant cross section variation was observed between the low frequency (a few dozen megahertz) and the high frequency runs, thus what we show here are aggregated statistics at V_{nom} and room temperature. This also means that SETs can be neglected and therefore, for our experiment–simulation comparison we just need to average the simulation results obtained under the eight static states. Note also that for this study, TIARA is configured with the behavioral BA-Ba collection model, which is well suited for SEU studies, as was elaborated in Section 3.2.2 on page 86. As can be seen in Figure 4.9, the simulations are about 1.5× pessimistic in the high-LET regime, and are compatible with the measured LET threshold: no errors were ever observed at 1.83 MeV·cm^2/mg, meaning that typical flip-flops have higher Q_{crit} than high-density memory cells, because of their larger transistor sizes. Overall, this validation bench demonstrates that from calibration against simple cells, namely SRAMs, we are able to extrapolate projections on more elaborate circuits such as this ca. 40-transistor flip-flop, with a confidence level fully in the range of accuracy that can be expected from Monte Carlo SER simulation.

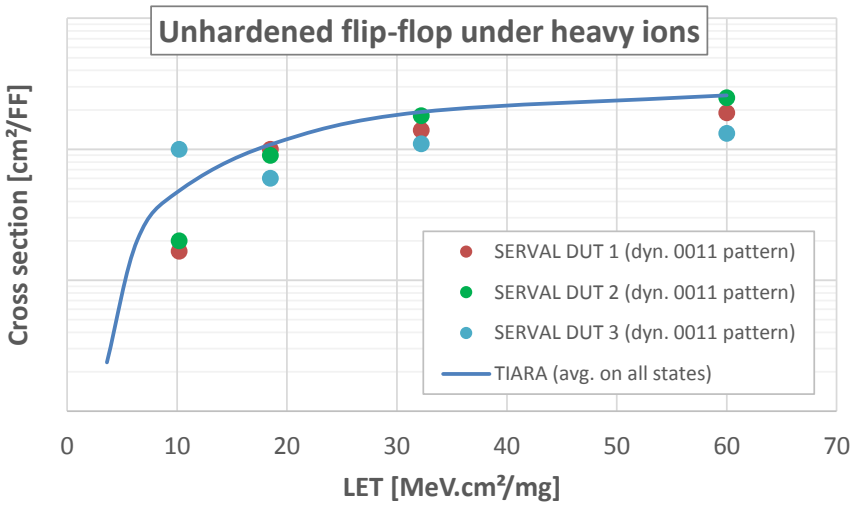

Figure 4.9: Model-Hardware Correlation (MHC) on SERVAL DFF under heavy ions.

More interestingly, in Figure 4.10 we show simulation results for high-energy protons, alongside experimental results with very large error bars: for the sake of completeness, many different DFF flavors were embedded in SERVAL, leading to shifters of relatively modest size for each type of cell. The shifter considered here is 20 kb, and because

protons are less aggressive than heavy ions, during the PSI campaign no more than one event per run was observed, leading to extremely poor error statistics and results that were hard to exploit. Fortunately, simulations provide us with more quantitative results regarding the proton response of this flip-flop, with cross sections being on average 5× below that of SRAM cells thanks to transistor sizing leading to higher Q_{crit}. Moreover, we can easily break down the sensitivity between the various electrical states, as shown with the eight dashed curves in Figure 4.10; according to the simulations, there is a 3× gap from the hardest to the softest state at 200 MeV, something that could not be disentangled from the experimental results.

All in all, this perfectly illustrates one of the major benefits of our Monte Carlo modeling approach, which was praised as early as Chapter 1: all that simulation takes

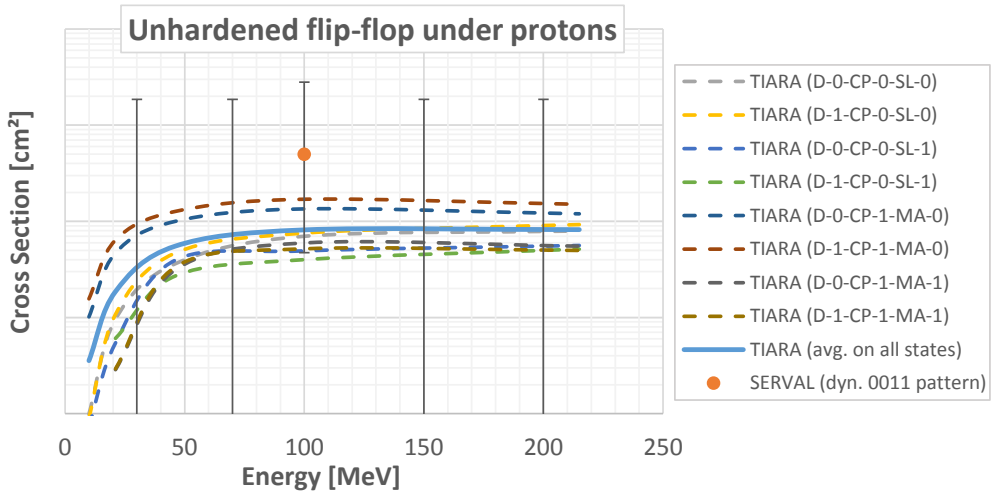

(a)

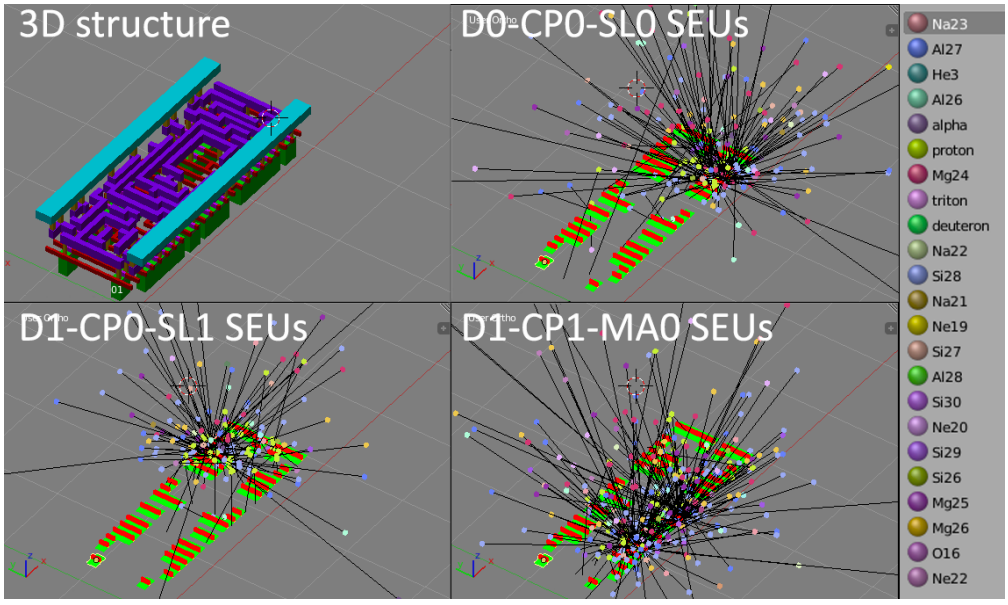

(b)

Figure 4.10: Proton projections on SERVAL flip-flop: (a) Simulated cross-section for each electrical state and experimental data with error bars at 95% confidence level and 10% dosimetry uncertainty – (b) 3D visualization of SEU-inducing recoils at 100 MeV.

is CPU time—something not specific to radiation effects—while beam testing activities incur significant costs in order to yield the *real* behavior of our circuits to radiation. However in the FD-SOI context we present here, standard accelerated fluxes are sometimes not sufficient to truly qualify circuits with ultra-low intrinsic error rates. In this example, our simulations give us precious characterization elements beyond the observability of normative, space agency-defined test protocols. Note that in a subsequent test chip, a different strategy was chosen in order to tackle this low-statistics issue, by selecting a more restricted catalog of cells. This allowed to inflate the shifters and collect satisfying results in the experiments, that proved to be in line with our expectations.

4.2 Single-event transients in planar bulk and FD-SOI processes

This section is devoted to the study of single-event transients, with a focus on 65 nm bulk Si and 28 nm FD-SOI. SETs have become increasingly studied these last years, starting from the 2000s; as mentioned in Chapter 1, this is mainly because the operating frequency of digital circuits has reached a point where glitches in combinational logic paths are quite likely to be latched by downstream sequential elements such as flip-flops. Therefore, the low temporal masking of SETs in modern CMOS nodes is such that SEUs are not necessarily the predominant contributors to the overall SER. Furthermore in the space industry's frame of reference, several efficient hardening solutions have been proposed against SEUs throughout the years; using ECC on memory arrays can eliminate Single-Bit Upsets (SBUs) and MCUs, provided proper multiplexing factors are chosen. In 28 nm FD-SOI, with less than 1% of memory events being MCUs [75]—and in most cases, two-cell upsets, Single-Error Correction, Double-Error Detection (SECDED) can harden memories against the vast majority of events. As for sequential logic, although the DICE architecture becomes less and less efficient in bulk with technology downscaling due to charge sharing mechanisms (see for instance [71]), recent measurements in FD-SOI technology [79] point out that the historical DICE design can still lead to flip-flops or latches that are radiation immune—or as close as can be, i.e. only with residual error sources from gazing-angle impacts or low-probability pulse propagation between the master and slave latches [144].

In brief, the current state of the art in RHBD allows to tackle single-event upsets very efficiently, at varying degrees of robustness in fact: based on catalogs of cells that can range from full heavy-ion immunity to softer designs addressing ionizing products from protons (or neutrons and alphas on Earth) with lesser performance impact, selective hardening can be used so as to surgically protect the critical elements in an SoC. In order to rationally quantify the criticality at cell level in complex circuits such as microprocessors, digital fault-injection approaches, most recently taking into account the circuit layout, have been demonstrated with very successful results [145–147]. Quite logically, SETs have thus been one of the topics put in the spotlight by the radiation community lately. They have been extensively characterized [148, 149], and impact analyses have also been carried out [74] along with propositions for mitigation (for instance with the so-called "pulse quenching" strategy [150]). We wish to address these last two topics here, first with advanced TIARA applications in 28 nm FD-SOI (Section 4.2.1), then by presenting experimental and simulation results on clock trees in 65 nm bulk (Section 4.2.2). Note that some content in Section 4.2.1 is taken from [14], and Section 4.2.2 is adapted from [16].

4.2.1 Explorations on combinational gates in 28 nm UTBB SOI

On the importance of radiation–circuit coupling: a worked example with transistor stacking

Throughout Chapter 3, we have claimed that the feedback loop between radiation-induced charge transport, and SPICE-level circuit response, is of particular importance for simulating SEEs in advanced technology nodes: as the voltage drops due to radiation charge injection, further collection is reduced because the electric field extinguishes. This sort of moderation law is not captured in our basic FD-SOI collection model BA-Ba, but constitutes one of the major features of our advanced response model CaTHODE, being a kind of custom 1D mixed-mode TCAD solver. In this section, we would like to show a worked example highlighting the relevance of circuit coupling, by considering a simple 28 nm FD-SOI RHBD inverter drawn in Figure 4.11a, with transistors stacked in series; device stacking is a well-known method for hardening, whereby transistors are connected with a common gate, and with their drain–source terminals in series, achieving the same logical function as a single transistor of identical connectivity. When a particle impinges on the circuit, it has to shunt both stacked devices for the output voltage to drop significantly, hence in SOI technology stacking is very effective since multi-device impacts are very rare [151–153].

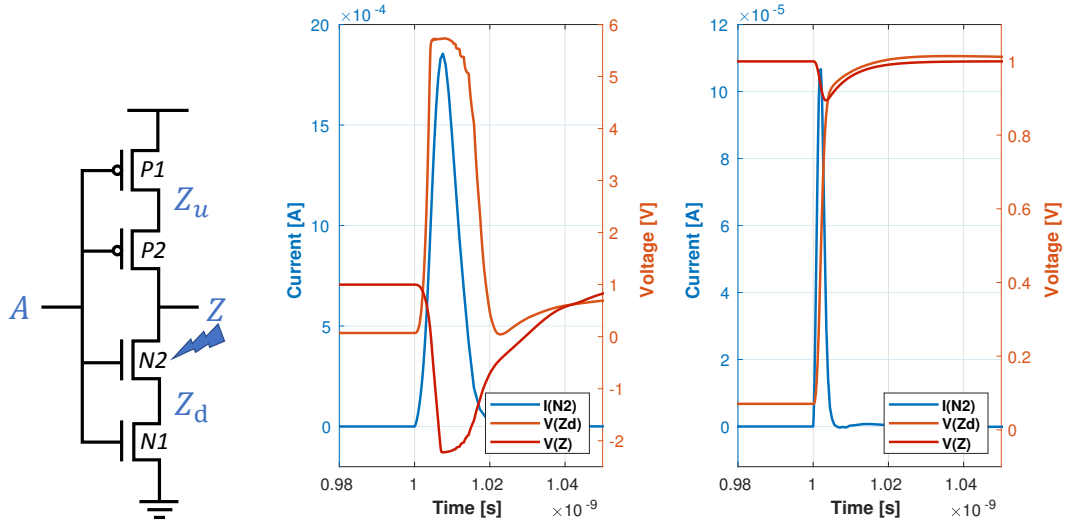

(a) (b)

Figure 4.11: Rad-hard inverter test case: (a) Schematic with stacked transistors – (b) Waveforms for 60 MeV·cm^2/mg impact on N2, without (left) and with (right) CaTHODE circuit coupling.

More rigorously, the robustness comes from the fact that the radiation-induced current vanishes with the impacted FET's drain–source voltage, and in return V_{ds} tends to clamp at 0 V rather than change sign for a prolonged time. This is illustrated in Figure 4.11b, where we show a vertical 60 MeV·cm^2/mg impact on transistor N2, simulated with CaTHODE for input pin A at logical '0', both with and without circuit coupling. The decoupled simulation is performed by omitting the voltage updates that the circuit solver normally feeds back to the CaTHODE structures. In other words, the transient 1D device simulation runs with V_l, V_r from (3.17), page 90, fixed at their values after quasi-stationary ramp-up to the injection-time target. This is equivalent to

the simulation scheme we apply to run our bias-independent collection models, in which we compute the radiation-induced currents assuming fixed terminal voltages, and then inject these frozen current waveforms into a separate electrical simulation afterwards. As can be seen from Figure 4.11b, in the decoupled simulation the voltage on output node Z drops largely below 0 V, while that on internal node Zd largely exceeds V_{dd}. On the other hand with the circuit-coupled simulation, we see that the radiation current is quickly dampened, avoiding unphysical values for the voltages of Z and Zd: their waveforms barely even cross as V_{ds} of nMOS N2 cannot change sign for long. Note that Zd's voltage rises more sharply that Z's potential drops, because Zd is floating, N1 and N2 being in 'off' state.

As a result of these radiation–circuit dynamic properties, only low-amplitude and short-lived SETs can be generated at output Z of the stacked inverter. Additional heavy-ion simulations were performed, where we simulated a short chain of inverters following the victim inverter and recorded SETs of magnitude above $V_{\mathrm{dd}}/2$. For an unhardened inverter (i.e. only transistors N1 and P1 directly connected to Z), there was no massive cross-section difference between the coupled and the decoupled simulation (20% increase), in spite of different SET widths predictions. As already justified in Section 3.2.2 on page 86 for simple collector BA-Ba, this is because on circuit portions without floating nodes (i.e. all our case studies up to this point), using a bias-independent current source the *occurrence* of digital SETs is correctly predicted, although with some conservatism. On the contrary, the stacked inverter was predicted to be immune up to 60 MeV$\cdot$cm^2/mg normally-incident ions, while without circuit feedback its cross section was comparable to a standard inverter's: with circuit coupling activated, the very faint generated glitches are almost always filtered by the next logic stages. This truly establishes the relevance of developing bias-dependent single-event models if one is to try and predict the robustness of a given circuit design: without circuit coupling, the simulator simply misses the entire point of hardening by stacking in SOI. Note that this hardening scheme, although very attractive from a robustness perspective, comes with non-negligible cost in terms of PPA: the switching speed drops by about 2$\times$ (since the saturation current $I_{\mathrm{on}} \propto W/L$ is about halved for the transistor pairs in series with $L_{\mathrm{eff}} \sim 2L$), and an area penalty of ca. 2$\times$ is also to be expected, depending on layout rules. In comparison and as will be discussed later on, more fine-grain control can be achieved by using transistors stacked in *parallel*, or multi-finger devices, with improved cell delay and good robustness (both due the larger I_{on}, that is doubled for $W_{\mathrm{eff}} \sim 2W$), at comparable area overhead.

Advanced TIARA case study: poly-bias usage in standard cell libraries

Carrying on with our SET investigations on 28 nm SOI combinational logic gates, in this section we discuss the influence of a particular design flavor called Poly-Bias (PB). Poly-biasing is a CAD-to-mask operation that retargets the layout-drawn polysilicon length to a slightly different physical length, allowing circuit designers to derive several standard cells from a single master and achieve a wide panel of speed and power consumption targets. In TIARA we account for poly-bias in the exact same way as described above, which gives rise to transistors with a new gate length L; here we will focus on the standard 28 nm FD-SOI inverter with $L_{\mathrm{min}} = 30$ nm or PB0, and its derivatives with 4 nm and 8 nm poly-bias (PB4 and PB8), respectively reaching $L = 34$ nm and $L = 38$ nm (all lengths are pre-shrink).

In Figure 4.12 we plot the collected charge Q_{coll} versus impact position for verti-

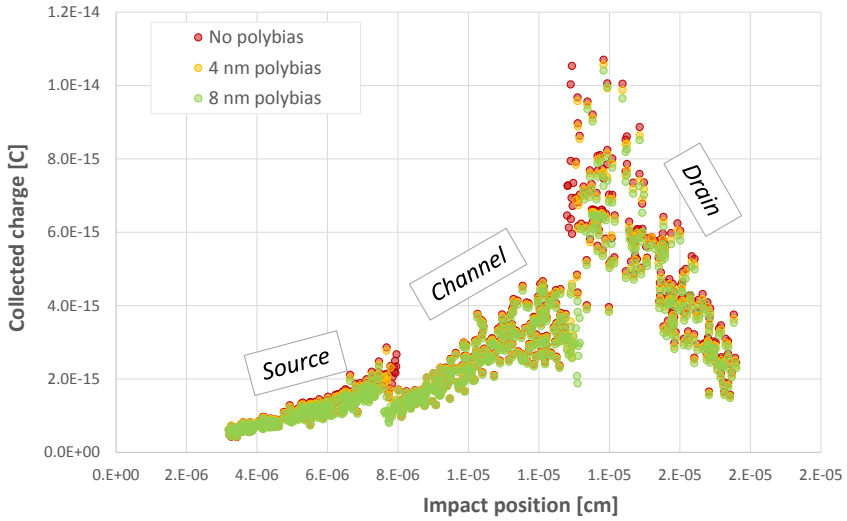

Figure 4.12: Charge collected at 60 MeV·cm^2/mg for several poly-bias variants of an inverter.

cal strikes on the various inverters' nMOS at 60 MeV·cm^2/mg. As can be seen, the maximum Q_{coll} is a slightly-decreasing function of PB, because the electric field is less intense in long-channel transistors. Note that statistical fluctuations on Q_{coll} are due to straggling in deposited charge Q_{dep} that we model for thin-film SOI (see Appendix B), hence Q_{coll} is not a single-valued function of strike location. Now in terms of overall response, this reduced charge collection of PBx inverters compared to PB0 leads to a small decrease in SET cross section (3% at most), as shown in Figure 4.13. Together with electrical characterization of the different PB flavors, this sort of graph can help orient the choice towards a particular design option.

From a radiation-effects point of view, our motivation for examining these poly-bias variations was that there are competing mechanisms at work in the overall response. These are summarized in Table 4.1 for poly-bias and another design option, namely the V_{t} flavor—RVT or LVT in ST's 28 nm FD-SOI. Be it for charge collection or circuit response, LVT is favored over RVT, while both flavors have equal area, and indeed, measurements on SERVAL test vehicle (not disclosed here) do confirm that LVT cells exhibit a somewhat lower cross section than RVT cells. On the other hand, for poly-bias we could not know beforehand what the global outcome was: PB0's higher Q_{coll} could very well have been outweighed by its higher Q_{crit} owed to a better drive current, or also by its slightly smaller sensitive area. Naturally, these results may also depend on the cell's intrinsic properties as well as its surroundings (the input driving signals and the output load), offering multiple axes for cell-level explorations.

Physical characteristic	**V_{t} flavor comparison**	**Poly-bias comparison**
Charge collection ($\mathcal{F} \implies Q_{\text{coll}}$)	RVT > LVT (steeper doping, stronger field)	PB0 > PBx (shorter transistors, stronger field)
Sensitive area ($W \times L \implies \sigma_{\text{sat}}$ in SOI)	RVT = LVT (same drawn devices)	PB0 < PBx (smaller channels, sensitive volumes)
Circuit response ($I_{\text{on}} \propto W/L \implies Q_{\text{crit}}, LET_{\text{th}}$)	RVT < LVT (slower transistors)	PB0 > PBx (larger drive)
Overall radiation trend	RVT usually more sensitive than LVT	? (not straightforward analytically)

Table 4.1: Concurrent radiation mechanisms for FD-SOI standard cell design options.

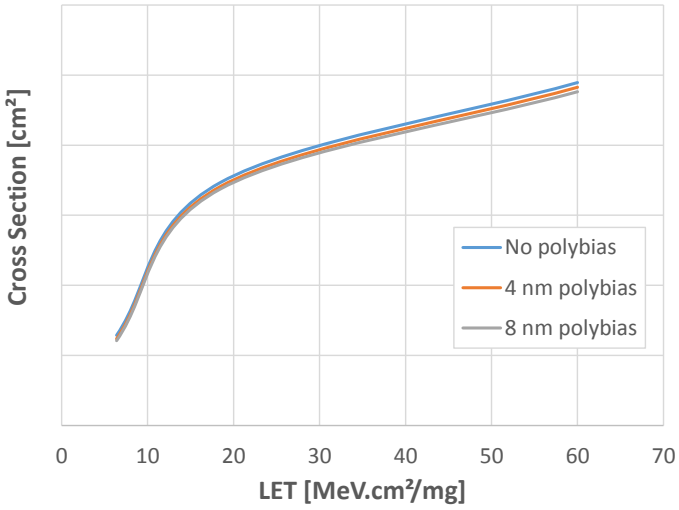

Figure 4.13: Heavy-ion cross section for poly-biased inverters (SETs taller than $V_{dd}/2$ after four chained inverters).

4.2.2 Pulse generation and propagation in 65 nm bulk Si clock trees

Having characterized the cell-level SEU and SET behavior of FD-SOI technology in depth, we now present a system-oriented discussion of SETs. Out of a criticality argument, we direct our study towards clock trees in digital SoCs, this time focusing on bulk silicon at the 65 nm node. Note that the presented results are in ST's 65 nm space platform [81, 82] built around specific process and DRC rules (e.g. to prevent SEL), robust cell libraries making use of redundancy, and rad-hard IPs qualified under beam; nevertheless, the cells we discuss hereafter are unhardened, and the trends we highlight could also apply to vanilla bulk silicon technology.

Context: on the criticality of asynchronous, high fan-out nets

Of all the subparts in an electronic system that are vulnerable to SETs, the clock tree is arguably one of the most critical, for several reasons: due to the sole amplification function of its main cells (usually buffers when excluding clock gating, multiplexing...), it does not exhibit logical masking. Neither does the clock tree benefit from temporal masking, unlike combinational data paths where glitches remain silent when generated outside the latching window of end flip-flops—note that this could be said of other asynchronous signals such as the reset or scan-enable commands. On top of that, topologies like clock trees have a high fan-out by nature, and the numerous cells that they drive behave erroneously all at once under such events. Thus the topic has been receiving a growing interest from the radiation reliability community over the years; in [154], the authors first assessed the criticality of clock errors in the global chip-level SER, then experimental evidence came in works such as [155], and various hardening schemes [156,157] and sensitivity analyses [158] were proposed. Here in this section, we report heavy ion test results on 65 nm shift registers in which large error counts were pinned on clock events, leading to cross sections only dependent on the clock buffers. We then present TIARA Monte Carlo simulations that achieve very good agreement with measured data, allowing us to perform further studies on the parameters driving the radiation sensitivity of clock networks.

Experimental measurements on shift registers – buffer cells cross sections

We have already insisted on how direct measurement of SETs is challenging: it takes high-speed IO pads and large sampling rates, or delicate custom structures such as VDLs to capture glitches with widths in the picosecond range. What was not mentioned yet, is the fact that exhaustive characterization of transients affecting combinational logic cells is virtually impossible. There may be only a handful of memory compilers in a technology offer, but standard cell libraries typically contain dozens to hundreds of gates. And while bit-cells see a constant load when they are in retention (only during read–write operations are they connected to the word- and bit-lines), logic paths are rarely evenly balanced. They can come in a variety of depths and fan-outs, giving rise to an exponential number of input slopes and output loads to browse through. Moreover, characterizing the likeliness of pulse *generation* can be insufficient; glitch *transmission* is also extremely sensitive to the propagation path, due to broadening or narrowing effects that strongly depend on the cells and on PVT [159]. Recent studies do try to tackle the above challenges however; in [160], an extensive characterization of transients was presented in 65 nm bulk at cell level using VDLs. The same authors then demonstrated a bottom-up SER estimation approach [161] in which the cell-level cross sections are aggregated to derive cross sections for more complex blocks. They use SPICE simulation to fill the gaps and estimate the masking effects. Because measurement points are not available for all combinations of cells, input signals, and output loads, several interpolating or extrapolating assumptions have to be made—still the agreement achieved by the authors is remarkable, despite imprecisions on the threshold-LET window.

The approach we choose here is different: we use shift registers that are normally dedicated to SEU measurements, and by classifying large error clusters as SETs stemming from their clock trees, we get an *in situ* characterization of pulse generation and transmission. Although very indirect, this method probes the real context in which logic gates are used, which was perhaps the most challenging in bottom-up approaches mentioned earlier. More specifically, a test vehicle called KIPSAT (Figure 4.14a) was designed and manufactured in STMicroelectronics' 65 nm space bulk silicon technology. It embeds SRAM arrays, microprocessors, microcontrollers, sensors, and shift registers which are the focus of this study. The shifter runs were performed at nominal supply voltage V_{dd} =1.2 V and at room temperature. Similarly to Figure 4.8 in Section 4.1.2 (see page 107), on-chip pattern generation and error counting is ensured by a BIST on each shifter, allowing for high operating frequencies under irradiation (225 to 600 MHz). The scan chain sizes were 13.5 kb or 28 kb, and the injected pattern was a checkerboard, i.e. 0, 1, 0, 1... Figure 4.14b illustrates one such shifter with an SET propagating in its clock tree, of reduced depth and fan-out for clarity. The heavy ion testings were carried out at RADEF, in compliance with ESA/SCC basic specification No 25100. The ion species used were nitrogen, silicon, argon, iron and krypton ions at various tilt angles, covering a range of LETs from 1.9 to 60 MeV$\cdot$cm^2/mg. Each run involving the shifters achieved a fluence of 5×10^6 ions/cm^2 at least.

The methodology we use to discriminate the clock tree errors is illustrated in Figure 4.15a for one of the runs at 32.1 MeV$\cdot$cm^2/mg (krypton). It shows data collected for two of the shifters, one of them made of standard flip-flops and the other made with robust flip-flops, but both clocked by trees made of the same buffer. The error counter is reset by the FPGA at each read cycle (every 0.75 s) to prevent overflow, thus the "error count versus time" signal (Figure 4.15a) is not monotonic. The robust flip-flop is more resilient than the standard one, which explains why the "background noise" is lower on the

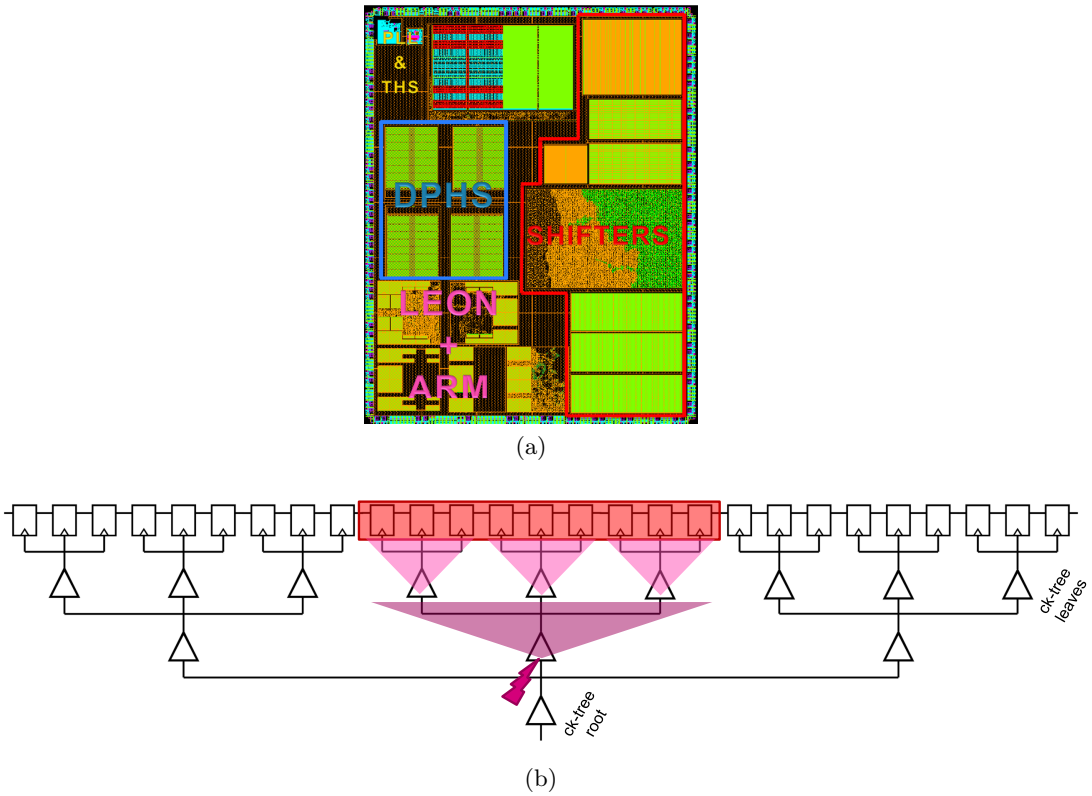

(a)

(b)

Figure 4.14: Overview of the studied structures: (a) KIPSAT test chip layout with shift registers on the right – (b) Schematic of minimalistic shifter with clock tree.

green curve. However, high peaks on the green curve are more frequent, albeit smaller, suggesting that these large error clusters originate in the shifters' respective clock trees. More quantitatively, the histograms of Figure 4.15b show that on average between two consecutive reads, the standard shifter sees about 18 errors when the robust shifter sees 12; this means that the robust flip-flops are about 3× harder than the standard ones (≈ 18/12×28 kb/13.5 kb). But the large events on the high-capacity shifter are more probable, although their size is generally around 60, against ≈100 for the smaller shifter. A quick analysis of their respective clock trees can explain these error sizes: the 28 kb shifter is clocked by a tree of depth 5 and whose internal fan-out is 8, and with the last-stage buffers driving 7 flip-flops each (i.e. $28 \times 1024 = 7 \times 8^{(5-1)}$). This means that clock-tree events originating in the second-to-last stage and above will induce shifter errors of multiplicity at least 7×8=56. The smaller 13.5 kb shifter on the other hand has a tree depth of 4 with an internal fan-out of 12 and a leaf-buffers/flip-flops fan-out of 8: $13.5 \times 1024 = 8 \times 12^{(4-1)}$—and thus the clock events from the second-to-last stage and beyond should have a minimum size of $12 \times 8 = 96$. This means that experimentally, "large flip-flop event" thresholds ($Th_{\text{large FF evt}}$) can be set to those values in order to discriminate the events coming from the clock tree excluding its leaves:

$$Th_{\text{large FF evt}} = FO_{\text{tree}} \times FO_{\text{to flops}} \tag{4.1}$$

where FO stands for fan-out. Below those thresholds, multiple events cannot be traced back to the clock tree with absolute certainty, since the flip-flop shifters are also prone to multiple-cell upsets in bulk technology. In other words, because intrinsic flip-flop MCUs in the registers and SETs in the tree leaves have similar error signatures, we can only

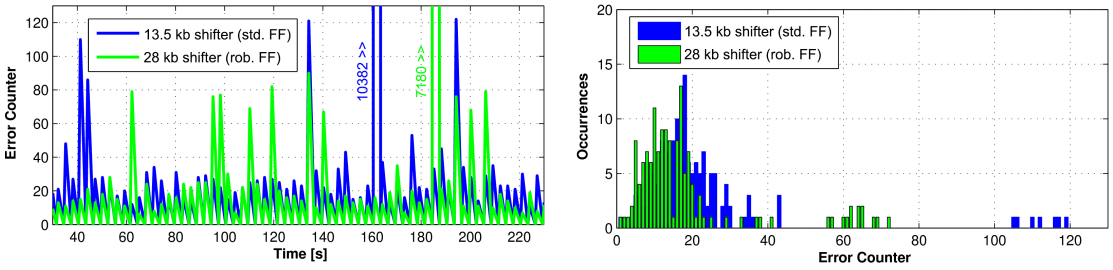

(a) (b)

Figure 4.15: Experimental heavy-ion data with (a) the raw error bus signal for a krypton run on two shifters, showing several peaks due to flip-flop error clusters from clock-tree SETs – (b) The resulting histograms of error count between two reads, with "large errors" standing out.

pinpoint SETs from the second-to-last tree stage down to the root. Note also how the highest peaks in Figure 4.15a indicate that, although rare, impacts closer to the tree root are indeed recorded and have the ability to upset large portions of a shifter all at once.

Using this criterion, we define σ_{bf} the event cross section of the clock buffers by counting $N_{\text{large FF evt}}$ the large flip-flop events on each run, and dividing by the fluence Φ times N_{bf} the number of buffers (excluding tree leaves):

$$\sigma_{\text{bf, exp}} = \frac{N_{\text{large FF evt}}}{\Phi_{\text{exp}} \cdot \underbrace{(N_{\text{bf, tot}} - N_{\text{bf, leaf}})}_{\sum_{j=0}^{d-2} FO^j}} \tag{4.2}$$

(where d denotes the tree depth, and dropping the "tree" subscript in FO for brevity from here on). These experimental cross sections are displayed in Figure 4.16 for all the shifters featuring various capacities and flip-flops. They prove to be essentially flip-flop independent, because the shifter clock trees use identical buffers. Apart from statistical uncertainty—clock events are quite seldom, as was seen in Figure 4.15—and from tilt angle effects, fluctuations are attributed to variations in flip-flop characteristics (minimum pulse width, input pin capacitance) and different tree topologies (depth, fan-out). Such experimental curves are extremely valuable, since they allow direct calculation of the clock-induced SER on other designs that use this buffer. More generally, shifters appear to be a quite direct way of capturing SETs due to their "transient-to-static" conversion ability. Clock trees might be challenging radiation-wise, but it is precisely the absence of temporal and logical masking that enables extracting meaningful data. Moreover, the regularity of our shifters makes their clock trees simple enough for analysis (no clock gating, constant fan-out...), yet real-life case studies.

Monte Carlo simulation results – clock-tree cells profiling

To deepen our analysis, we now make use of TIARA simulations monitoring the SETs affecting clock tree cells and leading to SEUs in flip-flops. The `Collector` module is set to use the so-called "Diffusion–Collection" model [7], which is our standard for bulk silicon simulation [5] while industrial integration of Random-Walk Drift–Diffusion (RWDD) in TIARA is still ongoing. Even though the model does not support dynamic circuit coupling—the injected currents only depend on node voltages calculated without

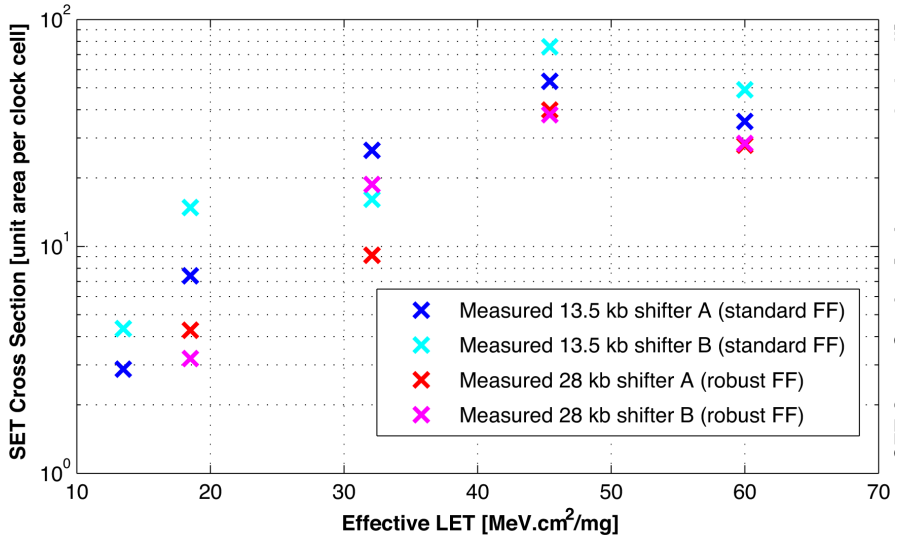

Figure 4.16: Experimental heavy-ion SET cross sections of clock tree buffers, nearly independent from downstream flip-flops.

injection—we can rightfully use it to predict digital SET *occurrence*: similarly to what was presented in Section 3.2.2 on page 86, the non-coupled model accurately determines whether the voltage drops, or if an SET happens or not. But because the injected current is not subsequently dampened, the model overestimates the duration of the voltage glitch, hence the SET width distributions we show later on should be considered pessimistic. Our simulated structures are made as close as possible to the aforementioned experiments: clock-tree buffers are instantiated with tunable depth and fan-out, and connected to a daisy chain of flip-flops. The simulated testing conditions are static: the tree root has constant voltage (0 V or V_{dd}), leading to a static pattern in the shifters (either 0, 1, 0, 1... or 1, 0, 1, 0...). To use the terminology of [154], this means we consider radiation-induced clock *race*, but not *jitter* (glitch simultaneous to a clock edge). This is a well-defined approximation given the ratio of SET widths ($\approx$ 100 ps – 1 ns, as shown later on) to the clock period (4.4 ns at 225 MHz for most of the experimental tests). To speed up the simulation of clock trees featuring a very large number of cells, we reduce the simulated netlist to a single tree path, as exhibited in Figure 4.17a for the same parameters as Figure 4.14b, i.e. $d = FO = 3$. Then at runtime, before drawing the cell-level impact coordinates, the simulator randomly picks one of the candidate cells with non-uniform probability across the stages; attributing a weight:

$$P(\text{stage i}) = FO^i / \sum_{j=0}^{d-1} FO^j \tag{4.3}$$

forces the impact distribution on the chain to what would be observed when irradiating the underlying tree. In other terms, each stage of the simulated path is given the statistical importance of its corresponding stage in the complete tree.

For each simulation, 100 000 impacts are computed, distributed across ten LET values, four possible electrical states (tree root and shifter state), and throughout the tree stages. All desired inputs–outputs are stored by TIARA's `Analyzer` module in a database on a per-impact basis (e.g. glitch amplitude and width as it propagates, resulting number of flip-flop SEUs if any, ...), allowing to display e.g. cell-level maps of impacts that gave rise to a latched glitch (Figure 4.17b), pulse width distributions (Fig-

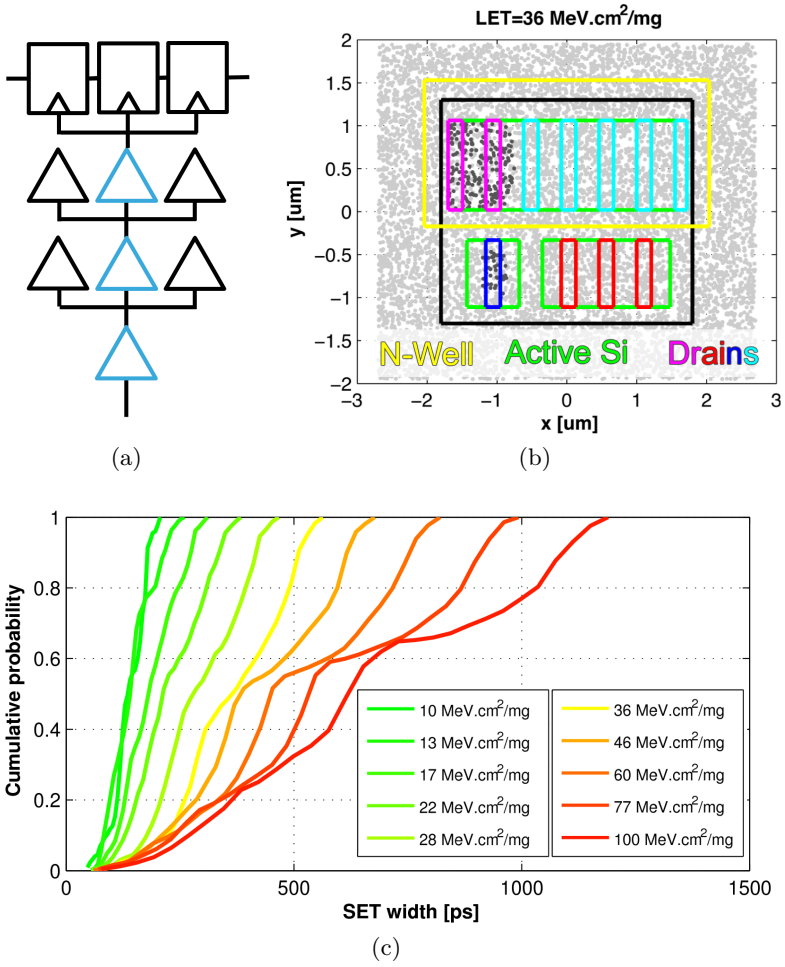

Figure 4.17: Example of clock-tree Monte Carlo heavy-ion simulation: (a) Irradiated tree path (impact candidates in blue, of impact probability scaled w.r.t. their population in the full tree via (4.3)) – (b) Cell-level "SET map" of the buffer (harmful impacts in dark dots, located at the off p-n junctions of blocked transistor drains) – (c) Pulse-width distributions in terms of Full Width at Half Maximum (FWHM), increasing with LET.

ure 4.17c)... To enable direct comparison with experimental results, again a buffer cross section can be defined from the number of flip-flop events; as we do not simulate irradiation on the flip-flops themselves, all flip-flop SEUs, regardless of the error multiplicity, originate from a clock-tree SET (i.e. a pulse that did not get electrically masked during propagation, and was wide enough to trigger at least one SEU among the flip-flops). Therefore, unlike in (4.2), page 116, glitches coming from the tree leaves need not be filtered out from the event count:

$$\sigma_{\text{bf, sim}} = \frac{N_{\text{FF evt}}}{\Phi_{\text{sim}} \cdot \underbrace{N_{\text{bf, tot}}}_{\sum_{j=0}^{d-1} FO^j}} \tag{4.4}$$

For comparison purposes, first of all a simulation was run with the same buffer and tree topology as the experimental 28 kb shifters (depth 5, fan-out 8). The latched-events cross section curve obtained is shown in Figure 4.18, in good agreement with experimental data. This validates the simulation methodology and justifies in-depth

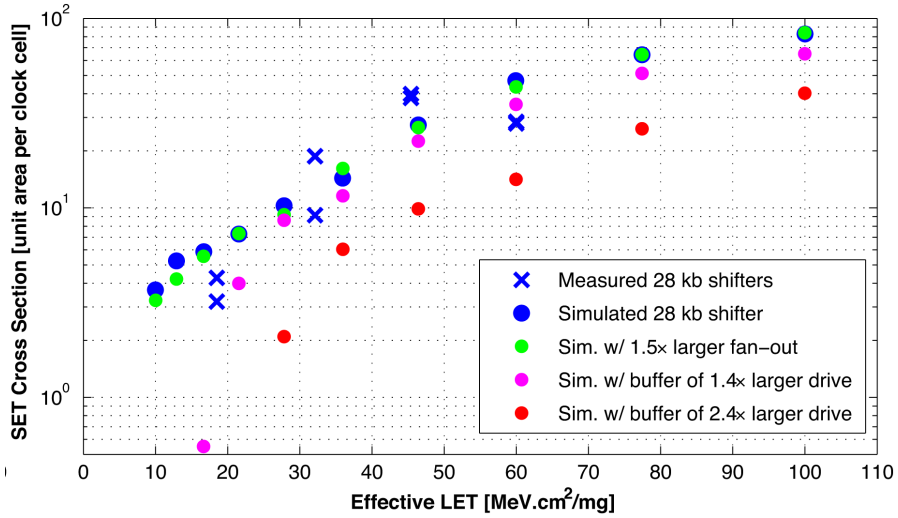

Figure 4.18: Simulated heavy-ion SET cross sections of clock tree buffers, for several configurations (identical as experiment, different tree shape, different buffer).

analysis of the simulation results.

One of the key outcomes in our simulations is that most of the clock SETs come from the tree leaves, as predicted in [154]. This is shown in Figure 4.19a where we plot a histogram of the tree level impacted for the latched glitches (with the simulation setup identical to the experimental 28 kb shifter, all LETs and input electrical states together). The event count is largely dominated by the last tree level, simply because the population of leaf buffers outweighs the populations of other levels. To further this analysis, in Figure 4.19b we plotted the histogram in terms of *flips* (total number of flip-flops SEUs) instead of *events* (oblivious to MCU sizes). If all tree stages were equally vulnerable, we would expect this histogram to be flat (and conversely, the bars in the histogram of Figure 4.19a would precisely scale by a factor of FO from one to another). However this is not the case, and the sensitivity decreases slightly towards the tree root, suggesting that some electrical masking occurs. Note that this was confirmed by further simulations targeting fixed tree levels: "bell-shaped" pulses can easily be captured by the flip-flops when they are generated at the leaf buffers, but closer to the tree root, they are more often flattened than sharpened into a digital-like glitch. Such considerations regarding vulnerability of the various tree stages are very valuable at system level, where the primary concern could be to estimate the likelihood of catastrophic clock events flipping thousands of flip-flops simultaneously.

To assess the radiation hardness of other clock trees, many more simulations were run on a variety of tree topologies and clock cells. We show a sample of them in Figure 4.18. Cell-level analyses are possible, e.g. oftentimes cells with stronger drive current are more resilient, as illustrated for buffers of similar design to the one irradiated experimentally, but with higher output drive current: a higher conductance means that given a certain radiation-induced current, the voltage swing at the cell output is smaller. Note however that the global outcome can be quite subtle, since increasing the drive current of a cell is usually achieved by adding transistors in parallel (multi-finger transistors). Thus from one drive to the other, the lowest cross section is not always achieved by the cell with higher drive current. This happens when the electrical benefit of a stronger drive current does not outweigh the area overhead of an increased number of fingers (yielding more

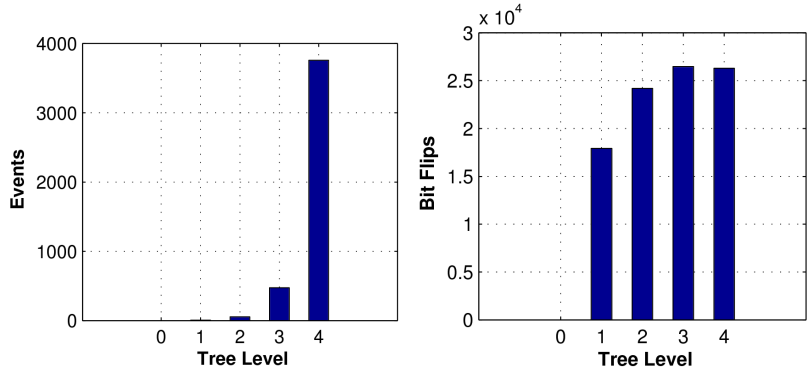

(a) (b)

Figure 4.19: Histograms of tree levels (numbered from root to leaves) giving rise to a captured pulse, (a) counting shifter errors regardless of their size, (b) counting total bit flips in the shifters. The tree leaves dominate the SET response.

sensitive areas for radiation-induced charge extraction)—this is observed on low-drive cells especially, whose layout varies more abruptly upon addition of transistor fingers. However the clock-tree cells presented here are quite massive compared to standard logic cells, since they generally have to drive ten cells at once, and that is why in this context, cells of higher drive current showed a reduction in cross section.

Aside from cell-level analyses, circuit-level studies were carried out, e.g. in Figure 4.18, the increased fan-out has weak effect. Note that on other cells that are not shown here, the fan-out can have a more pronounced influence, since heavier loading tends to flatten the SET pulses. In addition, the tree depth was also varied and no significant effect was seen, as expected from our previous observation that the tree leaves vastly dominate the error count. Finally, given the relatively small depth of the clock trees studied here, pulse broadening and narrowing, although observed in simulation (a few picoseconds per gate), are completely negligible. This is quite different from a "combinational logic" situation, routinely featuring paths of 10–20 gates.

Discussion: SET characterization methodology and mitigation rules

In this worked example of clock tree error analysis, we have shown that using a simple, but methodical filtering criterion for "large" multiple upsets in flip-flop shifters, we can trace back certain error signatures to glitches generated in the clock trees. This allows to extract experimental cross sections of clock cells nearly independent from downstream flip-flops. Although not as thorough as direct SET measurements, these cross section curves already provide quantitative information about the SET sensitivity of a given technology process, thereby extending the scope of radiation test chips that are usually focused on SEUs in SRAMs and flip-flops. TIARA Monte Carlo irradiation simulations have shown good agreement with the measurements, and have been used to perform analyses on untested structures, highlighting the influence of the tree topology (fan-out, depth...) or the cells it is made of. The context of clock trees proves to be very specific, and dedicated analyses are of utmost importance: unlike for SEUs, contextless SET simulations can only provide qualitative trends, and only through accurate simulation of the circuit specificities is it possible to deliver relevant predictions and hardening recommendations for designers.

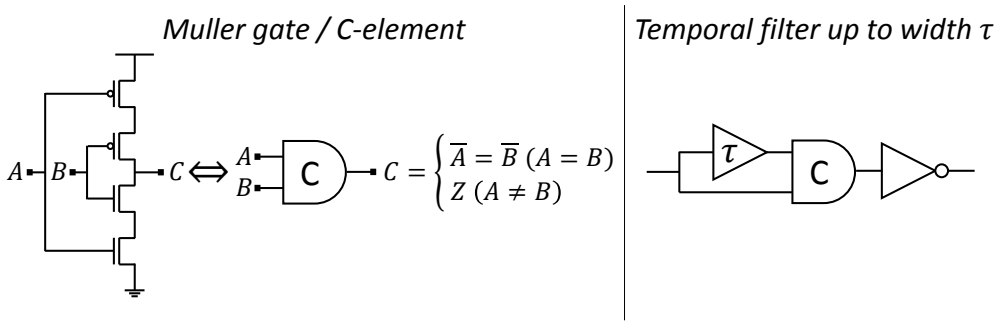

Figure 4.20: Buffer and C-element based SET filter: glitches narrower than delay τ are blocked.

In terms of SET mitigation rules, several options can be considered: to harden static nodes such as reset or scan trees, low-pass filters (e.g. realized with a passive RC circuit) offer the strongest safety levels. However, such analog solutions based on a Parametrized Cell (PCell) can be quite hard to implement for a clock tree generated in a stringent digital CAD flow. On the other hand, active temporal filtering solutions can conveniently be implemented in CMOS, as shown in Figure 4.20: a C-element or Muller gate is used to compare the signal to a delayed version of itself, which efficiently filters SETs that do not exceed the delay time [162]. In the time window during which the incoming SET is blocked though, the filtered output is at 'Z', or in high-impedance state. This makes it very sensitive to the surrounding radiating elements, and can cause signal integrity issues for such oscillating nodes as clock nets. Consequently, pragmatic SET recommendations for high fan-out nodes may be to try to hinder pulse *generation* rather than *transmission*. As shown by this study, this includes cell pruning based on transistor sizing, which can mean filtering below a certain output drive, or sometimes also inspecting the internal nodes of the cells.

4.3 On-orbit upset rate prediction

In this chapter's last section, soft-error rate projections in orbit are presented. Particle spectra are simulated for satellite trajectories around Earth, and the obtained fluxes are folded with energy- or LET-dependent cross sections for logic cells in order to yield an SER. In terms of modeling, the manner in which this flux–cross-section convolution is performed depends on several hypotheses; as we will discuss, TIARA allows to relax some of the assumptions in on-orbit SER models commonly used, while certain modeling choices still impose limitations on the types of environment we can natively handle. In terms of silicon technologies, we also quantify the SER benefit of fully-depleted processes compared to bulk, and discuss the industrial implications for trending space mission profiles. Note that some content in this section is adapted from [13].

4.3.1 Bulk Si and FD-SOI in geosynchronous Earth orbit

Introduction: on-orbit SER projection and environment models

Accurate prediction of the rate of single-event upsets is of utmost importance for space missions featuring a harsh radiation environment: while under-prediction of the system's SEU rate may result in mission failure, overestimating it can lead to an overly conservative design coming at a high cost, especially given the timescales involved. Because

direct measurements of SEUs on orbit—for technology qualification before production and launch—are quite rare, the most practical way to enable projections of a circuit's error rate in space is to use a combination of accelerated test data along with space environment models describing the particle spectra encountered on orbit. In many cases, the irradiation experiments can only be performed at normal incidence or just a few angles due to limited beam time; therefore, an extra modeling effort is required to extrapolate the response of the technologies under an omnidirectional space environment.

To that end, several SEU rate prediction methodologies have been developed over the years, for instance the Integral Rectangular ParallelePiped (IRPP) method originally developed in [60] and now implemented in popular tools such as CRÈME [18] and SPENVIS [19]. However, emerging technologies such as FD-SOI or FinFET feature strongly-confined geometries that are challenging the existing SEU models, as has been discussed throughout Chapter 3. Here we thus wish to discuss what modeling approaches are relevant in advanced depleted technologies. Note that in this section we focus on GEO environments, and hence we only consider direct ionization by the ion species on orbit, thereby neglecting upsets from secondary nuclear products; we do account for the trapped protons' contribution to the received LET spectrum, but overlook the upsets triggered by ionizing products of protons against silicon. LEO altitudes, i.e. 2000 km and below, will be discussed in Section 4.3.2, page 127.

In all of the following (including Section 4.3.2) we will use CRÈME simulations for our cosmic ray spectra. Similar fluxes can be obtained with SPENVIS and are not reported; as stated earlier in this manuscript, our goal is not to discuss the accuracy of radiation environment models, but rather to focus on the SEU models taking those fluxes as an input. Trapped protons are accounted for with the AP8 model, while Galactic Cosmic Ray (GCR) spectra are calculated with CREME96 GCR model. We cover multiple solar conditions, namely minimum and maximum "solar quiet" configurations (respectively maximizing and minimizing GCR fluxes) and "worst week" solar flare. All ion species up to Z=92 (uranium) are considered for flux calculations, and transported through variables thicknesses of aluminum. A sample of resulting LET spectra is plotted in Figure 4.21, where fluxes are given behind 100 mils spherical Al shielding[3], at a 0° longitude in GEO and 45° inclination in LEO. This figure truly illustrates the massive differences from one environment to the other, with LET curves being very distinct on a log–log plot. They differ not only in total flux magnitude—the two GEOs under different solar conditions are nearly homothetic above 2 MeV$\cdot$cm^2/mg—but also in LET content: the GEOs contains heavier species, while the LEO is dominated by light particles.

Bulk and FD-SOI compared with the rectangular parallelepiped model

In Section 4.1, page 99, we presented experimental heavy-ion cross sections for a 28 nm SPHD bit-cell in bulk Si and FD-SOI. These measurements are now used as a basis for extrapolation to real space environments: first-order predictions of the SEU rate are obtained via the Integral Rectangular ParallelePiped (IRPP) method, and we examine the validity of some crucial aspects in the model. Then based on the calibration procedure shown in Section 4.1 for the BA-Ba collection model, TIARA simulations allow us to derive a best estimate for the FD-SOI error rate in our case study.

The IRPP method [60], due to its genericity and large case coverage, has been widely adopted in the radiation effects community to compute SEU rates in space from ground-

[3]One mil is a thousandth of an inch, i.e. 0.254 mm or 25.4 µm.

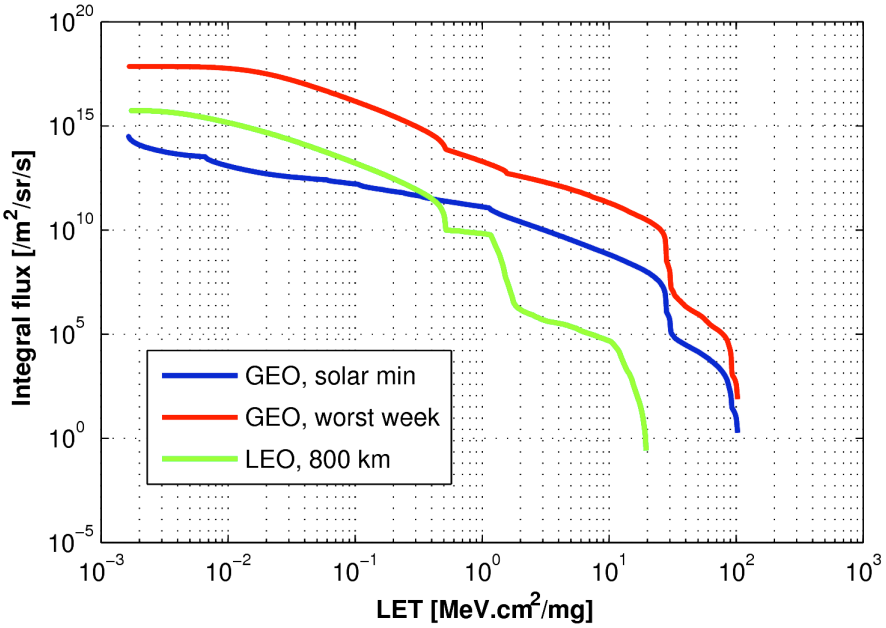

Figure 4.21: Integral LET spectra simulated with CRÈME in various environments: GEO under quiet Sun features high-LET species, while LEO or active Sun environments are dominated by low-LET protons and alphas.

based test data. The technical obstacle is to tackle practical calculation of the cross-section–flux convolution in (1.13), page 31: since σ can hardly be measured for all possible Z, E, θ, φ (the incoming ions' atomic number, energy, tilt and roll angles), several assumptions have to be made. First of all, the ion track structure (dependent on Z, E) is collapsed to the LET parameter, and then given a fixed critical charge Q_{crit}, the (non integral) RPP method works by considering a sensitive volume of x, y, z dimensions chosen to match both the experimental saturation cross section at normal incidence: $x \cdot y = \sigma_{\text{sat}}$ and threshold LET: $LET_{\text{th}} \cdot z = Q_{\text{crit}}$. Then the angular response is extrapolated by assuming that the sensitive volume has uniform 100% collection efficacy: for an ion of given LET and direction, this means that the collected charge will equal the charge deposited within the sensitive volume, yielding $Q_{\text{coll}} = Q_{\text{dep}} = LET \cdot l$ where l is the chord length traveled by the ion within the sensitive volume. This collected charge gives rise to an upset when it is superior to Q_{crit}. The probability of this happening is:

$$P(Q_{\text{coll}} > Q_{\text{crit}}) = \int_l f_l(l) \int_{L\,=\,Q_{\text{crit}}/l}^{\infty} f_L(L)\mathrm{d}L\mathrm{d}l$$

where—in this paragraph only—L is short for LET, and f_X is the Probability Density Function (PDF) for random variable X: $f_L = 1/\phi_{\text{tot}} \cdot \partial\phi/\partial L$ is the normalized *differential* LET spectrum, and f_l is the differential chord distribution inside the RPP. The above can be simplified to:

$$P(Q_{\text{coll}} > Q_{\text{crit}}) = \int_{L\,=\,LET_{\text{th}}}^{\infty} f_l(Q_{\text{crit}}/L)\frac{F_L(L)}{L^2} \cdot \mathrm{d}L \tag{4.5}$$

where $F_L(L)$ is now the Cumulative Density Function (CDF) for the LET, i.e. the *integral* spectrum ϕ/ϕ_{tot}. The SER is then obtained by multiplying (4.5) by $\phi_{\text{tot}} \cdot A/4$ where $A/4$ is the RPP's average projected area (see the discussion in Section 2.3.2,

page 59 for mathematical justification). In other words, the number of errors per unit time is the number of particles incident on the RPP per unit time, multiplied by the probability for *one* particle to cause an upset. Note that the $1/L^2$ factor in (4.5) simply arises from a variable change with $l = Q_{\text{crit}}/L \implies \mathrm{d}l = -Q_{\text{crit}}/L^2 \cdot \mathrm{d}L$, and it has an important signification: given the fact that for typical orbits, the flux drops rapidly with LET, this factor means that it is typically better to trade a low saturation cross section for a better threshold LET: when using wider transistors, σ_{sat} may increase (since in SOI it is driven by $W \times L$), but so will LET_{th} (driven by $I_{\text{on}} \propto W/L$), and the overall SER will be reduced. In early works, figures of merit have been proposed based on this $1/LET^2$ factor to reflect this notion [163].

In the RPP method derived above, the predicted cross section at normal incidence is necessarily a step function. Now, the more general IRPP method takes the above ideas further, by interpreting the $\sigma(LET)/\sigma_{\text{sat}}$ relative cross-section curve as a cumulated distribution function for the threshold LETs of RPPs of identical dimensions: while a step-function cross section curve corresponds to a single threshold LET, a gradual Weibull curve starting with an onset LET represents a wider distribution of threshold LETs from the onset LET value onwards. Note however that the physical meaning of a statistical superposition of identical sensitive volumes with different threshold LETs is quite questionable: at the time in 1992, it was believed that each individual bit-cell would exhibit a sharp cross-section curve, and that the smooth curve measured was due to a spread across the memory array due to process variability. We now know that even one cell alone intrinsically exhibits a gradual activation, which is the motivation for using nested sensitive volumes in certain Monte Carlo codes such as MRED [164].

Given an experimental cross section curve, the x and y dimensions of the IRPP can reasonably be estimated, starting from $\sqrt{\sigma_{\text{sat}}}$ (the default setting in CRÈME) and then playing with various aspect ratios. However, the choice of the z dimension can be more arbitrary; in bulk Si technology, it represents the depth over which quick collection by electric drift occurs, i.e. the depletion depth. The usual recommendation in CRÈME is to set z around 1 µm in planar bulk. For SOI technologies where charges deposited inside or below the BOX are not collected, the sensitive thickness can be much smaller than one micron; in our 28 nm FD-SOI technology, as discussed in Chapter 3, the silicon film is 7 nm-thin in the transistor channels and in addition, two factors can alter the notion of "effective" collection depth. First of all, parasitic bipolar amplification, which is typically below three for SEE-relevant LETs, increases the effective z. Second of all, detailed transistor morphology also plays an important role in defining sensitive areas, as was discussed earlier in this chapter; in our FD-SOI technology the active silicon thickness is not constant across the entire transistor (because of process features such as spacers, raised implants, silicide contacts), leading to 3× differences in silicon thickness outside the transistor channels. For this reason, a good estimate of silicon thickness in the region of interest (i.e. near the drain/channel interface) is 10 nm. For an in-depth analysis of the influence of elevated implants, see for instance [128].

Figure 4.22 depicts the z dependence of the IRPP SER predicted in the GEO environment at solar minimum shown earlier, for the 28 nm FD-SOI and bulk Si SRAM cross sections in Figure 4.2, page 101. While the bulk error rate—much larger than FD-SOI's—is relatively stable as z varies around its recommended value, the FD-SOI rate drops by a 30× factor from plausible z values to the z values recommended for bulk technologies. This suggests that very careful analysis of the IRPP outputs is necessary at advanced technology nodes with strong confinement at play. The increase in SEU rate with decreasing z, although counter-intuitive at first ("how can a smaller volume

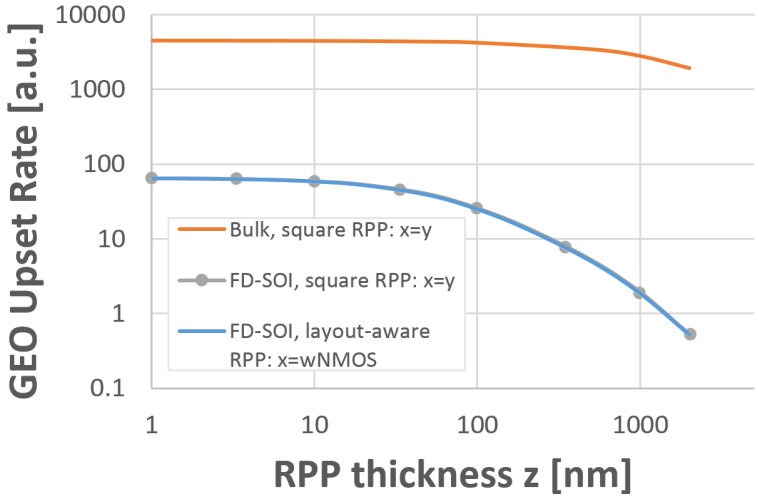

Figure 4.22: GEO solar min. upset rates obtained with CRÈME via the IRPP model for 28 nm FD-SOI and bulk Si SRAMs for various RPP dimensions.

be more sensitive?"), is easily understood either in geometrical terms (few chords are longer than z for a "needle", and the opposite for a "slab") or in electrical terms (Q_{crit} is low for a small z). These considerations also do apply to bulk technology, but in much lesser proportions than in FD-SOI: the sensitive volume in the bulk bit-cell being of micrometric lateral dimensions, when z is smaller than a micron the volume becomes a thin slab limited by z and the upset rate reaches a plateau: because the IRPP's Q_{crit} is scaled by z, in a slab of thickness z the portion of ion paths causing upsets is independent from z. In FD-SOI however, since the sensitive volume's lateral dimensions are an order of magnitude smaller than in bulk, when z is scanned across the range shown in Figure 4.22, it is either dominant or negligible, hence the large variations in upset rates. The IRPP x:y aspect ratio is investigated as well, using information from the cell layout: since the nMOS critical charge is much lower than for the pMOS and since they have much larger area in this SRAM design (unlike for combinational logic cells for instance), the sensitive volume x dimension can reasonably be approximated by the nMOS width. However, as shown in Figure 4.22, this has very little effect on the upset rate prediction, giving no more than 5% correction from the "square-IRPP" predictions. The difference becomes negligible as z gets much smaller than x and y, i.e. in the "slab" configuration where chord lengths are driven by z. This is why in Figure 4.22, the "layout-aware IRPP" curve is hidden behind the "square-IRPP" one. Such geometrical considerations can be very helpful to obtain orders of magnitude of the SER when foundry information is not directly available. For instance, note how a conservative estimate of the upset rate can be obtained by reading the "small-z plateau" in Figure 4.22, which is of practical interest when little is known about the detailed technology fabrication process. As an indication of the absolute numbers, the simulated upset rate for FD-SOI with z=10 nm is below 5 SEU/Gb/day, or 50× lower than the bulk Si SER with $z = 1$ µm.

Direct FD-SOI SER computation in TIARA

Now, using our dedicated response models for FD-SOI (in this case BA-Ba, which is sufficient for SEU simulation), a best estimate of the SER can be computed using all foundry information available for tight physical calibration. By setting up a TIARA simulation with an "on-orbit ions" irradiation scenario (Section 2.3.2, page 59), the SRAM

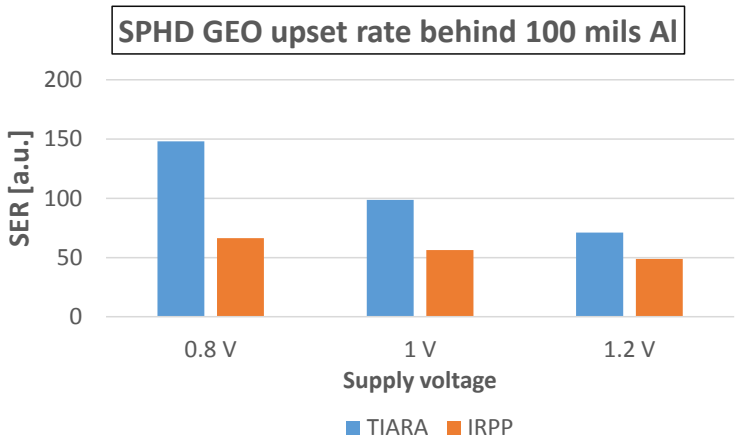

Figure 4.23: GEO upset rates obtained with TIARA for 28 nm FD-SOI SRAM, at different supply voltages and compared to IRPP estimates.

upset rate is calculated for our 28 nm FD-SOI bit-cell as the product of the integral flux by the omnidirectional cross section—under an isotropic, wide-spectrum flux. On a computational note, let us mention that some acceleration is needed for the simulations to run within a reasonable time: for sufficient statistics to be gathered, a few million impacts have to be simulated for each environment, which even with parallelization requires too much computing power. To avoid running all impacts entirely, we discard the calculations if the (initial) LET is below a certain cut-off, calculated to be the absolute minimum for an SEU to occur (by considering the longest chord in our sensitive volumes, the maximum bipolar amplification, the minimum critical charge...). For all environments, a huge portion of the LET spectrum stands below this cut-off, which provides a speed-up factor of about 100×, resulting in simulation times of about a few hours for one environment.

Direct comparisons between the IRPP upset rates and those calculated with TIARA in FD-SOI are given in Figure 4.23, in the same GEO environment (0° longitude, solar min., 100 mils Al) and at various supply voltages. For IRPP estimates, the sensitive thickness z is taken to be 10 nm as previously discussed, and voltage effects are *emulated* by assuming that the critical charge scales proportionally to voltage—thus modifying LET_{th} but not σ_{sat}. In this particular case the IRPP model underestimates the SER by 40–60% (if taking the TIARA simulations as reference). For practical purposes, this study shows that first-order estimations of the upset rates can be obtained with the IRPP algorithm when taking the silicon film thickness as IRPP z depth. This is mainly due to the fact that parasitic bipolar amplification remains modest at all times in FD-SOI; for technologies with a stronger parasitic gain (such as PD-SOI), this conclusion would probably not hold as the effective collection depth could be much larger than the physical silicon thickness above the BOX. Note however that the voltage dependence predicted by IRPP is not as steep as that predicted by TIARA, which is to be expected given the LET-dependent bipolar amplification that our custom response model accounts for. The full Monte Carlo simulations show a 1.5× increase in upset rate when the voltage is lowered by 20%. Conversely, strong reductions in SEU rate can be achieved by raising the supply voltage, offering interesting opportunities for hardened designs.

4.3.2 FD-SOI technology in low Earth orbit

Having discussed FD-SOI and planar bulk in GEO, we now move our focus to LEO environment. As developed in the first chapter of this dissertation (Section 1.5.2, page 41), the motivation for this is very timely: a host of academics and private companies are currently targeting short mission profiles in LEO, with low altitude enabling high-resolution observation or low-latency telecommunications. In that respect, it is important to characterize the soft-error rate of modern nodes such as 28 nm FD-SOI in LEO in order to support emerging space applications.

Modeling challenges for proton SER simulation

Compared to geostationary orbit, the main technical aspect to take into consideration for SEE modeling is the fact that the upset rate in LEO is often proton-dominated, due to the inner Van Allen belt (starting from approximately 1000 km) and to the magnetic shield efficiently deflecting GCRs (apart from high-inclination orbits). As far as modeling in TIARA is concerned, there are two main obstacles to direct proton SER calculation:

- Firstly, and as described in Section 2.3.5 on page 64, our nuclear databases were generated for normally-incident protons (and neutrons). This means that currently, integrating over all orientations to simulate an isotropic flux cannot be done. However, in Section 2.3.5 it was noted that overcoming this limitation is quite straightforward: nuclear recoil directions read from the input database simply need to be rotated depending on the primary particle, and to avoid biasing the simulation, one must account for the change in perceived thickness at tilted angles. Such developments are purely mathematical, and are currently ongoing in TIARA's `Irradiator` module. After they are completed, the platform will be able to simulate the SER due to high-energy protons (or neutrons) in arbitrary directions, and in particular isotropic protons for on-orbit environments.
- A more challenging aspect is that of Proton Direct Ionization (PDI) from low-energy protons, which has been shown to be a non-negligible contributing SEE mechanism for deep submicron technologies [165]. Although TIARA does include SRIM tables for proton stopping power, with the straight-ahead, Continuous Slowing Down Approximation (CSDA) that the `Raytracer` uses for charge deposition calculations, PDI-induced upsets are not properly captured: protons able to trigger an SEE are precisely those with an erratic trajectory, reaching sensitive volumes in the Bragg peak right before they come to rest. In Appendix B, we elaborate on how non-straight proton trajectories can be supported in future versions of TIARA.

Consequently, in the following, the proton upset rates we present for the 28 nm FD-SOI SPHD bit-cell are obtained under a simple isotropy assumption: σ is considered independent from θ, φ, and we only integrate over energy E to combine $\sigma(E)$ with the proton flux $\partial\phi/\partial E$. This is the same hypothesis used in CRÈME and SPENVIS for proton SER calculation. As shown in [166], for high-energy protons this assumption is reasonably justified, given the wide angular distribution of generated nuclear recoils; for low-energy protons, as discussed in [166] the SER can be mispredicted when using this isotropy assumption, especially in SOI where the sensitive volume's aspect ratio can induce strong orientation dependences. This certainly constitutes an axis of improvement to be researched. Given these approximations for proton-induced SER, in the following we make use of the simple IRPP model to compute the SER from heavier ions, again with $z = 10$ nm as sensitive depth.

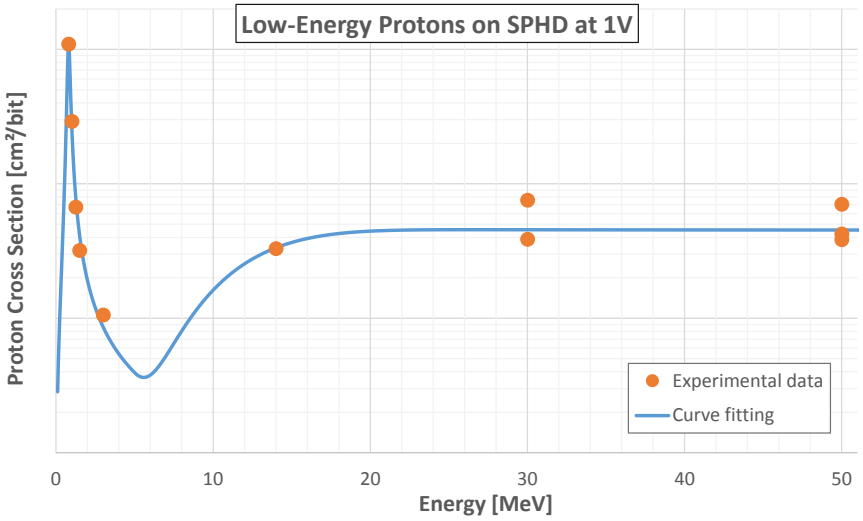

Figure 4.24: Experimental low-energy proton cross section for 28 nm FD-SOI bit-cell.

Influence of orbital parameters on the SER

In addition to the high-energy proton cross sections obtained at PSI and shown in Figure 4.6 on page 105, the 28 nm FD-SOI bit-cell was irradiated with low-energy protons at RADEF. The experimental beam results are shown in Figure 4.24. The low-energy cross section peaks at 0.8 MeV incident proton energy, where protons have a range of 11 µm, causing them to reach the sensitive volumes at their Bragg peak—the dice were delidded and irradiated from the front side, thus the thickness that photons need to traverse is that of the BEOL stack. In Figure 4.24 we reconstruct a continuous cross section curve by combining a standard 4-parameter Weibull for the high-energy regime, with a degree-two rational function to fit the data in the low-energy range.

Under the aforementioned assumptions for SER calculation, in Figure 4.25 we plot the upset rates projected for the FD-SOI SPHD bit-cell in various environments, for a shielding thickness of 100 mils or 2.54 mm, at solar minimum. As can be seen, the strongest driving parameter is the altitude, causing the error rate to increase between 0.5 and 1.5 decades every 500 km, since satellites enter the inner Van Allen belt from about 1000 km. The effect of orbit inclination is more complex: at 500 km below the radiation belt, moving towards the poles is detrimental (3× more SER per 25° inclination), firstly because of the South Atlantic Anomaly (SAA) and secondly, because the magnetosphere does not shield from GCRs—and for those polar orbits, even though they are LEOs, heavy ions can contribute a large fraction of the upset rate. On the other hand at 1500 km, higher latitudes are actually beneficial, due to the toroidal shape of the radiation belt (the SER is now proton-dominated), and the SER drops by 2× per 25°. At 1000 km finally, the SER is quite flat with respect to orbit inclination. For absolute numbers, note that the SERs shown in Figure 4.25 are below 100 SEU/Gb/day: the LEO upset rate can be higher than in GEO when inside the inner radiation belt, where numerous light particles cause more upsets than heavier but fewer cosmic rays for geostationary satellites.

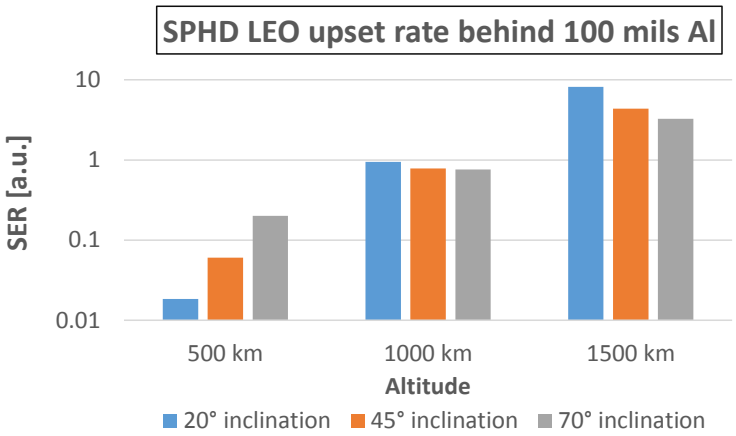

Figure 4.25: LEO upset rates projected for 28 nm FD-SOI SRAM for various orbital parameters.

Shielding and contribution of low-energy protons

Proton spectra on orbit are largely dominated by their low-energy component below 10 MeV. As shown in Figure 4.26 for a 1000 km/45° orbit with average space weather, a large portion of the low-energy flux is depleted thanks to shielding materials: it only takes 100 mils or 2.54 mm of aluminum (for instance a 1 mm-thin spacecraft hull and a 1 mm package for chips), to move the peak proton flux to about 30 MeV. Shielding thicknesses below this value that are drawn in the graph figure are representative of unshielded parts, i.e. with only a 1 mm package (40 mils) or a directly opened die (1 mil or 25.4 µm).

In this LEO environment, the respective SER contributions of low-, high-energy protons, and heavier ions, are displayed in Figure 4.27. As expected, for shielding thicknesses in the millimeter range, the upset rate is dominated by indirect proton ionization; on the other hand, for unshielded parts, the PDI contribution *to the proton SER* can be large (up to 75% for 1 mil). However for such thin shielding, the SER mostly arises from heavier ions, and overall for this orbit and space weather, low-energy protons only contribute 3% of the total SER at most. For other solar conditions, PDI-induced upsets may be more significant, for example in solar flare environments that are largely proton-dominated[4]. Nevertheless, looking beyond the mere goal of accurate SER prediction, for hardness assurance purpose, conventional methods for proton characterization are still largely relevant, as confirmed by this study in 28 nm UTBB SOI.

[4]Regarding direct proton ionization, let us mention that a larger discussion should also include dose effects—i.e. TID and TNID—in the picture, especially for components that are bound to be unshielded such as antennas or solar panels. It is important to understand that unlike single-event effects, as far as ionizing dose is concerned it is PDI that dominates, while indirect ionization via nuclear recoils is nearly negligible. A proton's maximum LET (about 0.5 MeV·cm^2/mg) may often not be enough to cause an SEE (which is why nuclear products at 5–15 MeV·cm^2/mg are mostly responsible for the error rate), but on average a lot more dose is deposited by protons directly than by their recoils: 0.5 MeV·cm^2/mg deposition every time generates much more TID than 5 MeV·cm^2/mg with a chance of 1/100 000–1/10 000 (as an approximation for the probability of nuclear interaction in relevant silicon volumes).

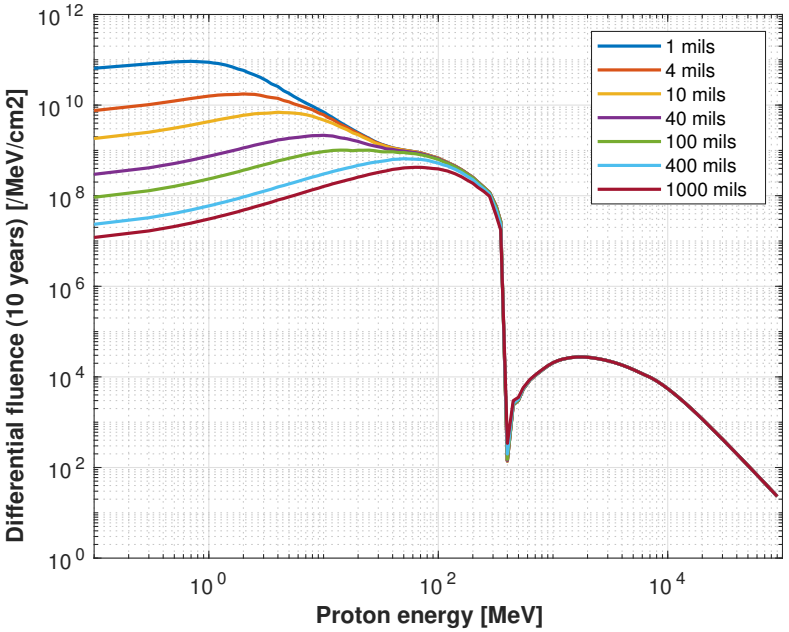

Figure 4.26: Ten-year averaged proton fluence at 1000 km and 45° at various Al thicknesses.

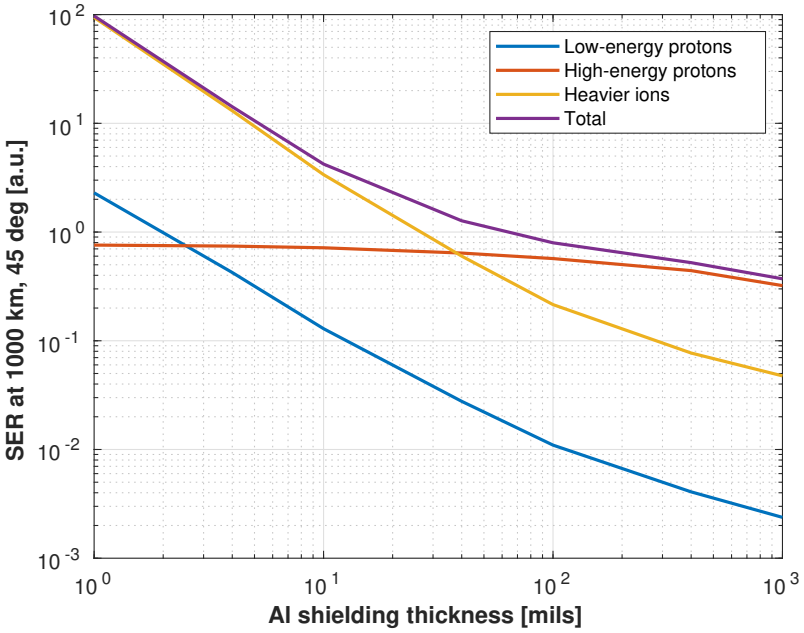

Figure 4.27: SER contributors for 28 nm FD-SOI SRAM in LEO.

4.4 Chapter conclusion

In this final chapter, we have discussed single-event effects simulation on CMOS logic cells, namely with SEU and SET studies in 28 nm FD-SOI and 65 nm bulk Si technologies, under various particle types and space mission profiles.

SEU measurements on SRAM cells with heavy ions and protons have provided direct evidence for the very high SEE resilience of FD-SOI technology compared to bulk Si. Furthermore, this collection of experimental data has allowed to calibrate the FD-SOI collection models developed in Chapter 3. After fine tuning both BA-Ba and CaTH-ODE against normal-incidence heavy ion data, simulations have shown to be predictive for projections in proton environments featuring a wide spectrum of nuclei, energies, and orientations. Based on this extensive Model-Hardware Correlation (MHC), TIARA has been used for in-depth characterizations of sequential and combinational logic cells in 28 nm UTBB SOI. With sequential logic studies, Monte Carlo SEE simulations have helped put precise figures where measurement gaps exist due to finite test condition sets. Then with standard and rad-hard inverter examples in FD-SOI, analyses of combinational cells were presented, highlighting the relevance of radiation–circuit coupling for SEE simulation at present-day process nodes, and allowing to explore the radiation robustness of several standard cell design options.

Subsequently, the system impact of SETs was discussed, with a study dedicated to clock trees in 65 nm bulk Si. Experimental cross-sections were extracted for clock buffers, and Monte Carlo simulations were run, yielding good agreement with the experimental data. Based on this, extrapolations to other cells and electrical contexts have been performed in order to explore hardening strategies. Within the timeline of this PhD, these SET studies have led to a patent grant for combinational logic hardening in bulk Si [167], and have given rise to practical SET mitigation guidelines in high fan-out nets such as clock or reset trees, both in bulk Si and FD-SOI. Indeed, such asynchronous signals are typically distributed to many, if not all, sub-blocks in an SoC, hence if they are left unhardened, they can be corrupted and lead to a SEFI. Throughout this PhD, these hardening guidelines were deployed and enforced both in R&D test chips and for customer products. The details of those rules cannot be fully disclosed here, however the general idea is that, since temporal and logical masking factors are very low in such trees, electrical masking of the SETs is the main lever for pragmatic mitigation. This is achieved by cell pruning in the library offer, in order to avoid the very generation of glitches, rather than their transmission.

Finally, on-orbit SER projections were operated in bulk Si and FD-SOI. Based on the IRPP model, both technologies were compared, and in geostationary orbit, 28 nm FD-SOI was found to bring an SER improvement between one and two decades compared to 28 nm bulk Si. Further heavy-ion simulations run with TIARA in FD-SOI were carried out, allowing to refine the SER estimates. The IRPP assessments were shown to remain within 0.5–2× accuracy when compared to Monte Carlo predictions with our custom FD-SOI response models, depending on the LET content of the defined orbit and space weather. 28 nm FD-SOI was also examined in LEO environments, especially for upsets induced by low-energy protons at low shielding thicknesses. In LEO outside solar flare conditions, it was shown that Proton Direct Ionization (PDI) was not preponderant in the overall upset rate, although the proton error rate of unshielded parts may be dominated by low-energy contributions.

Overall, this chapter has established how predictive SER simulation capabilities can

be used jointly with experimental investigations to practically address the emerging radiation hazards in modern ICs. As a concluding remark as to TIARA's range of applicability, note that the last section in this chapter was interspersed with hints at possible evolutions: currently, upsets induced by electrons and low-energy protons are simply "dismissed" in the simulator, and to estimate PDI-induced SER we had to entirely resort to normal-incidence experimental data. Taking into account these "subthreshold" mechanisms certainly is an area of future work.

Conclusion and perspectives

Cosmic rays, or energetic particles from outer space, have been known to affect the behavior of electronic circuits since the 1950s. When they are not shielded by the Earth's magnetosphere, space-bound microelectronic devices can present an intolerable rate of random failures, or SER. Because of this, historically radiation-hardened ICs for space have been based on conservative designs, and manufactured in process nodes lagging a few generations behind mainstream technologies in use for consumer applications, because of the several years needed to derive and qualify rad-hard platforms starting from commercial technologies. However, these practices are currently evolving, with many launches targeting LEO, serving commercial companies seeking to develop cost-efficient products within fast time-to-market. As a result, the space industry is increasingly turning towards highly-integrated CMOS technologies to serve its roadmap. In that respect, this thesis work has allowed to accompany the space industrial revolution, by improving the understanding and modeling of single-event effects of high-end CMOS technologies intended for space applications, especially focusing on UTBB FD-SOI.

Radiation environments and physical processes leading to single-event effects were reviewed in Chapter 1. The state of the art in soft-error simulation was also discussed, with TCAD tools serving device-level studies, and Monte Carlo simulators addressing SER projections at circuit level. Modern CMOS scaling trends have been discussed w.r.t. single-event upsets and transients, highlighting the increasing contribution of SETs to the overall SER, and showing the SER improvement natively brought by fully-depleted technologies compared to planar bulk transistor architectures.

Chapter 2 gave a detailed description of TIARA, the Monte Carlo tool used, maintained, and upgraded, throughout this thesis. The physical models implemented by TIARA were discussed, along with their limitations; in particular, the question of radiation–circuit coupling was addressed with care. Software-engineering aspects were treated as well, with a strong emphasis put on automation of the simulation flow allowed by automated 3D structure creation from the circuit layout, and its mapping into the electrical domain by an LVS-like procedure. Parallel execution was also discussed, and finally TIARA was benchmarked with other state-of-the-art tools.

Chapter 3 was devoted to the study of fundamental single-event mechanisms in FD-SOI. TCAD simulations on elementary transistors demonstrated the technology's native robustness, owed to small charge collection volumes and limited parasitic bipolar amplification. This physical insight was then used to establish two compact SEE models in FD-SOI; the first is mainly behavioral, and being circuit-decoupled, it is intended mainly for SEU simulation; the other model consists of a 1D transport simulation of the carriers generated in impacted transistors, coupled to the electrical solver in TIARA.

Thus it captures complex phenomena, such as parasitic bipolar amplification and dynamic circuit effects, from first semiconductor principles and in agreement with detailed TCAD simulations.

Finally, Chapter 4 was dedicated to SER simulation studies with TIARA. The two FD-SOI collection models developed in Chapter 3 were validated against beam test data collected on memory cells. TIARA was then used to study 28 nm FD-SOI sequential and combinational logic cells across several test cases, allowing to gain insight into the technology response beyond experimental observability. Furthermore, SER predictions were also discussed in 65 nm bulk Si, with a study of SET generation and propagation inside clock trees leading to practical recommendations for SET mitigation. Finally, the on-orbit SER behaviors of planar bulk and FD-SOI were compared, evidencing the benefits of FD-SOI technology for timely space programs.

As a result, the initial research objectives were fulfilled: the physical single-event models developed have proved to correlate with experimental data, and have allowed to complement silicon measurements: the SER of measured cells has been refined, that of unmeasured—sometimes even untestable—ones has been projected in orbit, and we have been able to draw upon these simulations to consolidate space platforms with pragmatic RHBD guidelines. This was done within a time frame shorter than typical silicon iterations, and compatible with emerging space programs.

As a reflection of the industrial nature of this PhD, perspectives to this thesis work are of two kinds: first from a scientific standpoint, improvements can be brought to the physics we model in our SER simulator; then in the engineering frame of reference, opportunities can be explored to optimize the hardening methodologies and practices.

First of all, possible improvements to the TIARA simulation platform have been discussed throughout this dissertation. Evolutions of the `Irradiator` module could include the integration of Geant4 databases to extend the range of materials that we simulate for nuclear interaction. Rare events stemming from recoil nuclei heavier than silicon, for instance, can be fundamental for proper simulation of rad-hard cells with a high upset threshold. As for direct ionization computed by the `Raytracer`, future work will be to account for detailed ion track structure, not only in terms of radial profile, but also in terms of direction straggling near stopping, which is essential to properly capture light-particle upsets. Finally, the `Collector` module could be extended to support FinFET technology; when developing CaTHODE, our 1D SEE model in FD-SOI, suggestions were given as to possible modeling strategies in FinFET. A hybrid modeling approach could be implemented to separate volumetric transport in the bulk, from confined transport in the fin.

Secondly and moving towards the upper abstraction levels, several bridges can be considered to enhance the radiation-hardening flow. Integration of SEE waveforms within the CAD flow is one of the keys to enabling rad-hard circuits on large scales, and enable the deployment of efficient hardening guidelines and practices. Although work in this area was not mentioned in this manuscript, embedding bias-dependent radiation current sources within a PDK is at least as equally important as having the most accurate error rate prediction: it is key to interactive feedback at circuit design time, considering the exponential number of cells, input slopes and output loads, or PVTs to be tested, which we cannot realistically hope to all address with Monte Carlo simulations. Similarly, synergies with digital fault injection tools, i.e. simulators operating at gate or RTL level, are a fascinating area of research to be explored. Topological, or layout-aware fault-injection was demonstrated only recently; by clever use of Analog and Mixed Signal

(AMS) simulation features, this could allow to generalize computational SEE analyses to large SoC scales not accessible to analog tools such as TIARA. Rather than hardening high fan-out nets that are *intuitively* deemed critical, systematically deriving *objective*, statistical criticality metrics could be a powerful decision aid for performing selective hardening with limited PPA penalties.

More generally, if within the next years the space industry roadmap does confirm to truly propel radiation hardening out of its niche market, with an ever-increasing number of chips to harden, bandwidth issues are to be expected. This is when sharing and training becomes crucial. In a way, TIARA or other radiation models such as TID-induced parametric drifts, are but a few among many CAD products. In that respect, radiation engineers and circuit designers have to tightly interact, similarly to the working model with EDA vendors. Clearly, seamless integration of Monte Carlo SER simulation in a standard CAD flow is yet to be achieved. But although it may not appear too much in this manuscript—which is necessarily science-oriented—with the TIARA developments carried out during this PhD, we have come a tremendous way towards industrializing the tool, and the road ahead truly is thrilling.

* * *

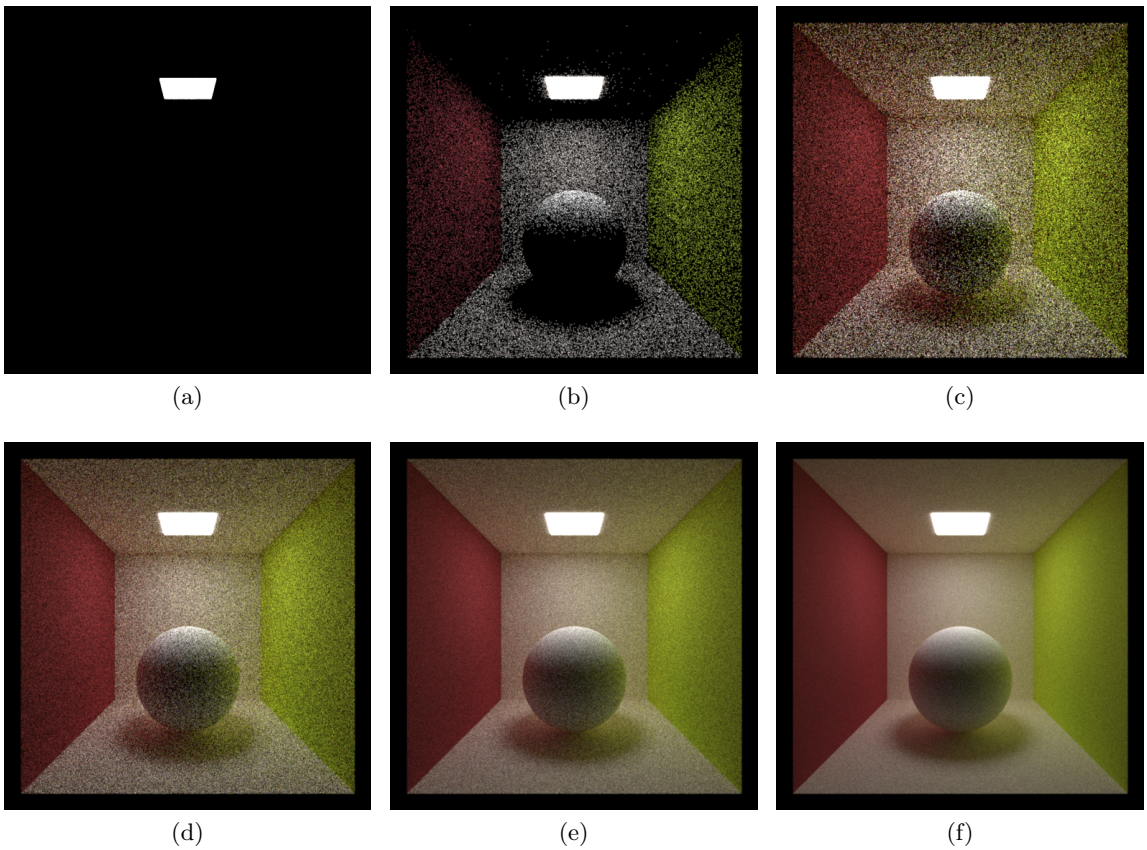

(a) (b) (c)

(d) (e) (f)

Photon transport simulation in TIARA (100-line Easter egg): (a) Max. scattering depth = 0 – (b) Max. scattering depth = 1 – (c) Unlimited scattering depth (d) Samples per pixel = 128 – (e) Samples per pixel = 512 – (f) Samples per pixel = 2048.

Appendix A

Theory and practice of stochastic sampling in TIARA's irradiator

In this appendix we briefly present Monte Carlo integration, the mathematical theory underlying TIARA's simulation scheme. We then detail the various sampling strategies implemented in its `Irradiator` module for generating particles based on spectral and geometrical parameters.

A.1 Theory of the Monte Carlo method

A.1.1 From deterministic to stochastic integration

In their general form, Monte Carlo algorithms are a set of methods that aim at evaluating the integral of a given function using random sampling. For the sake of simplicity here, let us consider a strictly positive function on interval $[a, b]$: $g : x \in [a, b] \mapsto g(x) > 0$, whose integral we wish to calculate on the interval. The simplest way of performing numerical integration is via a Riemann sum:

$$I = \int_a^b g(x)\mathrm{d}x \approx \frac{b-a}{N} \times \sum_{i=1}^{N} g(x_i) \tag{A.1}$$

where the x_i evaluation points form a regular grid over the integration segment. Non-uniform grids may be used to alleviate any issues around singularities of g, and higher-order approximations are often employed for better convergence properties—e.g. the trapezoidal rule uses a piecewise-linear approximation for g instead of piecewise-constant, yielding quadratic convergence instead of linear. But the main issue faced by all *deterministic* numerical integration methods is that their algorithmic complexity is polynomial with a degree given by the dimensionality of g and therefore, they are typically chosen for low-dimension problems. On the other hand, general problems arising from particle physics require infinite-dimension spaces since the number of particle interactions can be arbitrary. From a technical perspective, note how the sum needs to be entirely recomputed if a better approximation is desired: going from N to $N' > N$ displaces the entire evaluation grid, unless in special cases e.g. $N' = 2N$ where we refine

the grid by precisely halving the integration step.

The central point in Monte Carlo methods, which were developed in the 1940s precisely to tackle particle physics, is simply to consider a *random* evaluation grid instead. In our exposition example, let $X_i \sim \mathcal{U}[a,b]$ be a set of independent and identically distributed uniform random variables over the segment. Noticing in (A.1) that $(b-a) = \frac{1}{1/(b-a)} = 1/f_{X_i}$ is simply the reciprocal of the Probability Density Function (PDF) of X, we can derive a stochastic estimator for the integral from a size-N sample of X_i:

$$I_N = \frac{1}{N}\sum_{i=1}^{N}\frac{g(X_i)}{f_X} \tag{A.2}$$

where in fact the PDF f_X can be chosen arbitrarily. By construct, the expectation value of our Monte Carlo estimator is the integral of g:

$$E(I_N) = \frac{1}{N}\sum_{i=1}^{N} E\left(\frac{g(X_i)}{f_X}\right) = \frac{1}{N}\sum_{i=1}^{N}\int_a^b \frac{g(x)}{f_X(x)}\cdot f_X(x)\mathrm{d}x = \int_a^b g(x)\mathrm{d}x = I \tag{A.3}$$

In other terms, the estimator is unbiased and in practice, we can use one of its *realizations* to compute a numerical value for the integral. This is exactly what lies in the definition of a heavy-ion cross section at normal incidence for example; given N_{SEU} upsets under a fluence Φ of N impacts spread out over an $[a,b]\times[c,d]$ irradiation rectangle,

$$\sigma = N_{\mathrm{SEU}}/\Phi = \frac{(b-a)(d-c)}{N}\sum_{i=1} N\mathbb{1}_{\mathrm{SEU}}(x_i,y_i) \tag{A.4}$$

where $\mathbb{1}_{\mathrm{SEU}}(x,y)$ equals one if an impact at (x,y) triggers an SEU, and zero elsewhere. This definition of the cross section can in fact be thought of as an approximation to the integral of the *indicator function* of the sensitive surface: it is completely analogous to (A.1) generalized in two dimensions, with $g = \mathbb{1}_{\mathrm{SEU}}(x,y)$.

As a brief summary, Monte Carlo integration schemes are always composed of three bricks: sampling, evaluation, and aggregation. In TIARA, the `Irradiator` module is the sampler, in charge of generating the (x_i,y_i) in the above example; the `Raytracer`–`Collector`–`CircuitSolver` chain as a whole is the evaluation procedure, which sets the zeros from the ones in $\mathbb{1}_{\mathrm{SEU}}(x,y)$; finally the `Analyzer` module, which is responsible for the read–write operations on the event result database, also performs the aggregation step by summing the terms in the formula above. After a short parenthesis on general convergence properties of Monte Carlo integrators hereafter, the rest of this appendix will merely give the pragmatic details of the formulas used to implement the various irradiation scenarios allowed in TIARA.

A.1.2 On convergence properties and importance sampling

Back to our one-dimensional construction, if the X_i are still assumed independent, the variance of the random estimator is inversely proportional to sample size N: *the Monte Carlo method converges as* $1/\sqrt{N}$. This is somewhat of a slow asymptotic behavior, but insertion of new evaluation points is completely straightforward, unlike with deterministic algorithms. More formally, according to the central limit theorem, the estimator converges to a Gaussian distribution that narrows around the sought integral I:

$$I_N \underset{N\to\infty}{\sim} \mathcal{N}(I, \delta g/\sqrt{N}) \tag{A.5}$$

where:

$$\delta g = \sqrt{\int_a^b \frac{g^2(x)}{f_X(x)} \mathrm{d}x - I^2} \tag{A.6}$$

quantifies the dispersion of g measured under the chosen PDF f_X. Ideally, the variance of the estimator could thus be made null by drawing the samples with a PDF precisely proportional to g: $f_X = g/I \implies \delta g = 0$. This means that instead of computing the integral as a sum of $g(x)$ at uniformly-placed points, we could compute it as a sum of constants "evaluated" at points distributed with a g-shaped density. But this choice for f_X is ineffective since the constant I is the quantity to estimate. However, drawing the samples from a distribution that *resembles* the integrand is possible, provided we are able to normalize the chosen "shape" into a PDF. For instance, suppose we have forgotten the integral of the sine function between 0 and π but we still remember how to integrate piecewise-linear functions; we could generate samples from a triangle-shaped density mimicking a sine bell and perform the Monte Carlo integration, achieving faster convergence than with uniformly drawn samples.

The general method of "placing samples where it matters" (i.e. where the integrand is largest) is called *importance sampling*. In TIARA, this could be implemented by using *a priori* knowledge of which radiation events are more harmful. For example under isotropic irradiation, we could draw more samples at tilted angles and, for the estimation to remain unbiased, give said samples proportionally less weight when calculating aggregates such as the cross section. This is mostly an *implementation* challenge but poses no mathematical issue. Going even further, we could analyze the events already simulated in order to *a posteriori* alter the distribution for samples to come: as the cross-section map reveals, we could avoid resampling areas that are not sensitive, to focus on refining the SEU/no-SEU frontier. This would then make the samples correlated, which is more challenging *mathematically*. Finally, let us note that importance sampling preserves the average (expectation) values, but not the "physical fluctuations" (variances): if rare events are statistically enhanced, then the obtained set is no longer representative of an experimental data set where all events have equal weight. For all those reasons, in TIARA's `Irradiator` module the choice was made to only use natural distributions, perhaps at the cost of some convergence speed. As mentioned in Chapter 2, various pruning strategies are implemented to avoid processing unnecessary events all the way through irradiation, ray tracing, charge collection and circuit response. If pruning is efficient enough that rejected events take up negligible time with respect to a full processing sequence, it amounts to a particular case of importance sampling: quickly rejecting *a posteriori* useless samples is equivalent to extreme importance sampling where the PDF for sample generation has been dropped to zero in *a priori* insensitive places.

A.2 Techniques for sampling arbitrary distributions

To generate samples from an arbitrary probability law, the starting point is to at least know how to generate uniform random samples. That is the goal of a Pseudo-Random Number Generator (PRNG) such as the Mersenne Twister [96, 97], which we use in TIARA's `Irradiator`. This algorithm is overwhelmingly popular for Monte Carlo simulations. It was invented in the late 90s and soon replaced the linear congruence generators that were used almost ubiquitously at the time. This widespread adoption was mainly due to its much better statistical properties—e.g., how "random" the pseudo-random sequence is in terms of uncorrelated numbers—and was also owed to its huge period

(of $2^{19937} - 1$ by default)—while in old software packages, PRNGs with periods of just 2^{32} were common, which can turn out to be problematic. Another neat property of the Mersenne Twister is k-dimensional equidistribution, or the fact that all possible k-tuples occur an equal number of times over the period of the generator—which is a much stronger property than just having "well-distributed" points in k-dimensional space, i.e. without pattern. For instance in TIARA, this means that if we irradiate a $[0, 1] \times [0, 1]$ rectangle by generating (x, y) couples, not only will the points be satisfyingly distributed after a while, but in fact, eventually the whole grid will be evenly covered up to machine precision, meaning that absolutely all impact points will be inspected[1]. Shortcomings of the Mersenne Twister include the fact that it cannot generate parallel streams, and that it is cryptographically insecure. The latter is because only a few observations of the sequence are enough to determine where in the sequence we lie, and thus be able to predict all the subsequent values in spite of the very large period—but in TIARA this is obviously not an issue. The flaw of not providing independent parallel streams is caused by the fact that sequences starting with different seeds are not necessarily uncorrelated, especially if seeded with values close to each other. In TIARA this may be an issue to work around if we were to instantiate the `Irradiator` at slave level instead of having a unique PRNG on the master process (see Section 2.5.2, page 73). But to sum up, with our current software design choices the Mersenne Twister is a sound choice, and at present there is no need to look at alternatives for pseudo-random number generation.

In the rest of this appendix we will denote $U, u_1, u_2, u_3 \sim \mathcal{U}[0, 1]$ independent uniform random variables on $[0, 1]$. As we will see, when applied relevant transformations onto, they allow to generate any targeted distribution.

A.2.1 Rejection sampling from a probability density function

When given a certain density f_X, the most simple way of going about generating a sequence of numbers according to f_X is to use *rejection sampling*. Provided f_X has finite support $[a, b]$ and is bounded: $\exists M > 0, \forall x \in [a, b], f_X(x) \leq M$, then by:

- drawing uniform samples (x, y) on $[a, b] \times [0, M]$ from random variables $X = a + (b - a) \cdot u_1$ and $Y = M \cdot u_2$
- keeping the x values for (x, y) couples that fall below the $y = f_X(x)$ curve

we generate samples distributed according to f_X. This is illustrated in Fig A.1 and merely corresponds to a geometrical interpretation of the PDF as a continuous generalization of what a discrete histogram means. Rejection sampling is very straightforward to implement, but it becomes inefficient when the PDF has a narrow peak whilst having a large extent. In that case the bounding rectangle in Figure A.1 is largely unoccupied, and most of the samples are rejected.

A.2.2 Inverse transform sampling from a cumulative density function

A different approach is to use a Cumulative Density Function (CDF) for the law we wish to sample. Suppose we apply a bijective transformation g to uniform variable $U \in [0, 1]$,

[1] Another way of advocating for k-equidistribution goes like this: however unlikely it is to win the lottery, one would rightfully feel cheated if the winning ticket was never going to be sold anyway, even at very low odds. Or said differently, a k-equidistributed monkey sitting at a typewriter not only may, but eventually will, produce the works of Shakespeare.

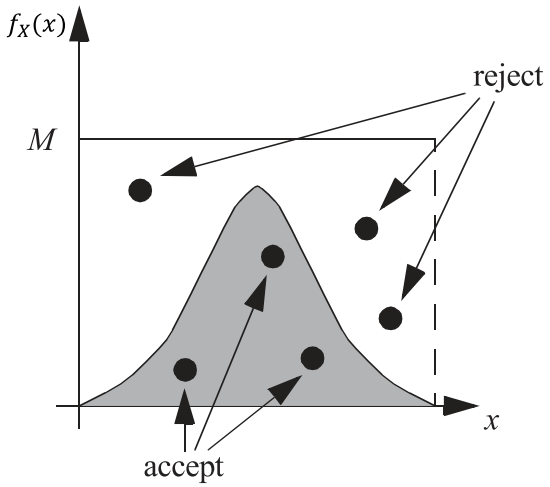

Figure A.1: Rejection sampling – adapted from [168].

to obtain variable $Y = g(U)$. Assuming g is a monotonically increasing function to dismiss sign issues, we can compute the CDF of Y by:

$$F_Y(y) = P(Y \leq y) = P(U \leq g^{-1}(y)) = F_U(g^{-1}(y)) = g^{-1}(y)$$

where we just used the fact that the CDF for U is identity on $[0, 1]$. Thus the CDF of $Y = g(U)$ is g^{-1}: applying transformation g on uniform variable $U \in [0, 1]$ distorts it into a variable of density $(g^{-1})' = 1/g'$.

Working out the above equation backwards, if we aim for a given f_X density, we see that:

$$X = F_X^{-1}(U) \tag{A.7}$$

where $F_X(x) = \int_{-\infty}^{x} f_X(t)\mathrm{d}t$, will generate appropriately distributed samples. In other terms and as depicted in Figure A.2: integrate the density, find the inverse, and apply it to uniform $[0, 1]$ numbers. Thus the method is called *inverse transform sampling*, and in the `Irradiator` this is what we use in the emitter's spectrum sampling method responsible for drawing particles according to a certain energy-dependent flux (see Section 2.3.2, page 59).

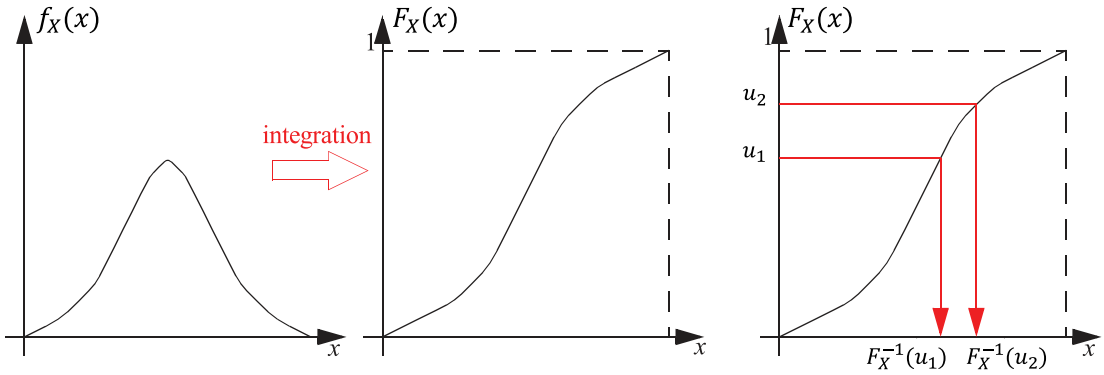

Figure A.2: Inverse transform sampling – adapted from [168]. Note that numerically, finding the inverse function is not a problem *per se* as it merely corresponds to a "reverse" interpolation from ordinate to abscissa.

To relate the on-orbit flux to a probability distribution, one just needs to normalize it by its integral value: if $\partial\phi/\partial E$ is the *differential* flux i.e. the number of particles per unit area per unit time per unit energy (with energy *between* E *and* $E + dE$), then

$$f_E = \frac{\partial\phi}{\partial E} \Big/ \int_0^\infty \frac{\partial\phi}{\partial E} \mathrm{d}E \tag{A.8}$$

is the PDF for generating random particle energies distributed according to the spectrum. Let aside the units and the mathematical formalism, this just means that the number of particles with a given energy is the total number of particles times the probability for *one* particle to have that given energy. Note that depending on the tool used for space environment simulation, sometimes we may work on the *integral* flux instead: $\phi(E) = \int_E^\infty \partial\phi/\partial E \cdot \mathrm{d}E$ is the number of particles per unit area per unit time with energy *below* E. With this right-tailed definition, the integral flux is a decreasing function of energy and signs are reversed compared to our increasing-CDF convention previously. Then the normalization factor i.e. the total flux over the whole spectrum is simply $\phi(0)$, and the integral flux is proportional to the CDF F_E for drawing energies: $F_E = \phi/\phi(0)$. The "integration" step is then skipped in the inverse transform sampling procedure.

A.3 Sampling volumes and surfaces for emission or impact points

A.3.1 Uniform points in a box volume

Drawing random points inside the volume of an axis-aligned box is the goal of the emitter's volume sampling method in the `Irradiator` (Section 2.3.4, page 62 and Section 2.3.5, page 64). To achieve a uniform distribution of points in an $[a, b] \times [c, d] \times [e, f]$ rectangular parallelepiped, we just need to compute their Cartesian coordinates as:

$$x = a + (b - a) \cdot u_1 \tag{A.9}$$
$$y = c + (d - c) \cdot u_2 \tag{A.10}$$
$$z = e + (f - e) \cdot u_3 \tag{A.11}$$

Or in other words, we draw independent uniform triplets and rescale them to match the appropriate volume, yielding a PDF of $f_{(x,y,z)} = 1/(b-a)(d-c)(f-e)$. This would also work in any other dimension.

A.3.2 Piecewise-uniform points on a box surface

Drawing random points on the surface of an axis-aligned box is the goal of the target's surface sampling method in the `Irradiator` (Section 2.3.2, page 59 and Section 2.3.3, page 61). Picking a point uniformly on one face is done exactly as above (Section A.3.1), expect in lower dimension—i.e., sampling inside a 2D rectangle instead of a 3D box. Then to achieve a facewise-constant density of points on the box surface, one first has to pick which face to sample out of the six possible choices. In the general case, we may wish to assign arbitrary statistical weights w_i to the faces ($i = 1 \dots 6$). Picking which face to sample is then a direct application of the inverse transform sampling procedure explained in Section A.2.2, for a discrete CDF: if we denote the normalized cumulated

weights $W_i = \sum_{j=1}^{i} w_j / \sum_{j=1}^{6} w_j$, then the index j of the chosen face is given by:

$$j = 1 + (\max i, W_i < U) \tag{A.12}$$

A.4 Sampling the unit sphere for outgoing or ingoing directions

In this section we will denote $\boldsymbol{u}$ a unit vector of Cartesian coordinates (u_x, u_y, u_z) and spherical coordinates (θ, φ) for tilt or zenith, and roll or azimuth angles. Generating propagation directions for our rays is equivalent to drawing points on the surface of the unit sphere and its subsets.

A.4.1 Uniform direction on the sphere

Generating uniform outgoing directions for particles is the task of the emitter's direction sampling method when the emitter directivity is set to be isotropic (Section 2.3.4, page 62). To distribute points uniformly on the unit sphere, one common mistake being made is to draw them from uniform spherical coordinates $(\theta, \varphi) \sim \mathcal{U}[0, \pi] \times [0, 2\pi[$. But doing so, the points obtained are bunched near the poles, the same way meridians are closer to each other around the poles than at the equator—or mathematically speaking, because the solid angle element $\mathrm{d}\omega = \sin\theta \mathrm{d}\theta \mathrm{d}\varphi$ is a function of θ.

A correct, easily implemented method to pick uniform points $\boldsymbol{u}$ *on the sphere* is to draw points $\boldsymbol{r}$ *in the sphere* via rejection sampling, then project them at a unity radius:

$$\text{try } \boldsymbol{r} = (x, y, z) = (2u_1 - 1, 2u_2 - 1, 2u_3 - 1) \tag{A.13}$$

$$\text{until } x^2 + y^2 + z^2 \leq 1, \tag{A.14}$$

$$\text{then } \boldsymbol{u} = \frac{\boldsymbol{r}}{||\boldsymbol{r}||} \tag{A.15}$$

This generates a properly uniform PDF $f_{\boldsymbol{u}}(\theta, \varphi) = 1/2\pi$. As illustrated in Figure A.3, at each try the rejection rate is the complement of the sphere-to-cube ratio of volumes, i.e. $\alpha = 1 - 4\pi/3 \cdot 1^3/2^3 = 1 - \pi/6 \approx 48\%$. Therefore, the average number of attempts needed is $1/\alpha \approx 2$, as the expectation value of a geometric distribution—the first time success of consecutive Bernoulli trials. Thus in spite of being of unpredictable execution time in theory, in practice this rejection method is quite efficient, and this is what we use in TIARA.

A.4.2 Uniform direction on a cone

In isotropic mode, to avoid generating rays that do not reach our targeted areas, the emitter's direction sampling function can also draw uniform directions in a cone (Section 2.3.4, page 62), or more rigorously a spherical cap of half opening angle $\theta_{\max}$ as depicted in Figure A.4a.

To achieve such a density of $f_{\boldsymbol{u}}(\theta, \varphi) = 1/2\pi(1 - \cos\theta_{\max})$ with $\theta \in [0, \theta_{\max}]$, we

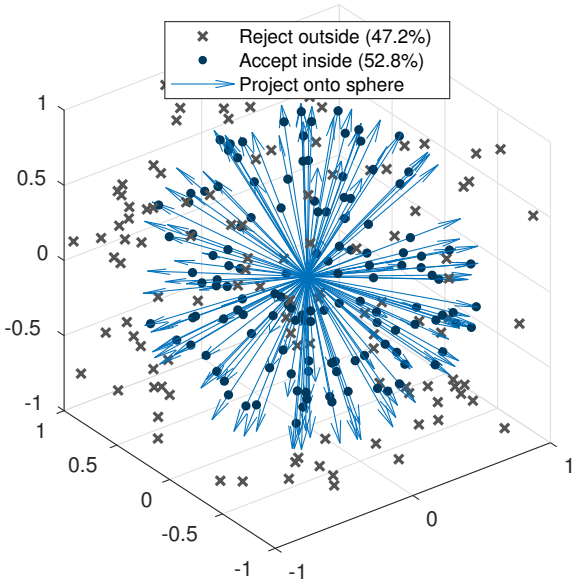

Figure A.3: Uniform sampling on the sphere.

compute:

$$u_x = \cos(2\pi u_1)\sqrt{1-[1-u_2(1-\cos\theta_{\max})]^2} \tag{A.16}$$

$$u_y = \sin(2\pi u_1)\sqrt{1-[1-u_2(1-\cos\theta_{\max})]^2} \tag{A.17}$$

$$u_z = 1-u_2(1-\cos\theta_{\max}) \tag{A.18}$$

In spherical coordinates, this corresponds to:

$$\theta = \arccos[1-u_2(1-\cos\theta_{\max})] \tag{A.19}$$

$$\varphi = 2\pi u_1 \tag{A.20}$$

This means that we choose a non-uniform zenith angle (what is uniform is the cosine of the inclination), and then a uniform longitude on the corresponding circle of latitude.

Note that this method can be used as an alternative to the rejection method presented in Section A.4.1, by setting $\theta_{\max} = \pi$. We then have $\theta = \arccos(1-2u_2) \sim \arccos(\mathcal{U}[-1,1])$, which is a direct application of inverse transform sampling: we are precisely inverting the sine-dependent distortion of the solid angle that was mentioned above: $|(\arccos^{-1})'| = \sin$. Which method works fastest certainly depends on computer hardware, whether the considered machine uses tabulated or Taylor-expanded trigonometry functions for instance. The rejection method bears the advantage of being easily generalized in arbitrary dimension—although the acceptance ratio falls off with dimensionality as the hypersphere occupies an ever-decreasing fraction of the hypercube. Finding analytical formulas for hyperspherical coordinates, on the other hand, is a bit more challenging.

A.4.3 Cosine-weighted direction on the hemisphere

Finally, when considering isotropic irradiation seen in the receiver's frame of reference, we need to draw ingoing particle directions with a probability proportional to the cosine of the local inclination (Section 2.3.2, page 59). This PDF of $f_{\boldsymbol{u}}(\theta, \varphi) = \cos\theta/\pi$ is implemented in the target's direction sampling routine. We first pick uniform points in the unit disk via rejection sampling, and then we project them onto the unit hemisphere:

$$\text{try } \boldsymbol{r_d} = (u_x, u_y) = (2u_1 - 1, 2u_2 - 1) \tag{A.21}$$

$$\text{until } u_x^2 + u_y^2 \leq 1, \tag{A.22}$$

$$\text{then } u_z = \sqrt{1 - ||\boldsymbol{r_d}||^2} \tag{A.23}$$

As can be seen in Figure A.4b, the projection stretches the points by a $1/\cos\theta$ factor, thus depleting the distribution towards the equator precisely by the right amount.

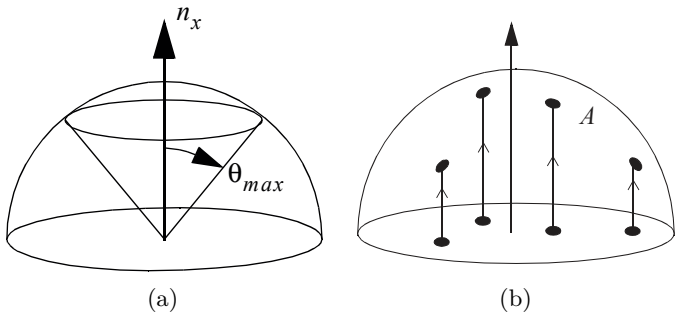

Figure A.4: (a) Uniform sampling on a spherical cap – (b) Cosine-weighted sampling on the hemisphere. Both taken from [168].

Appendix B

Ion track structure in TIARA: existing models and perspectives

This appendix discusses how the ion track structure is currently accounted for in TIARA, and how more elaborate models can be implemented in future evolutions of the tool. Although not uncorrelated, three aspects need to be considered beyond the straight-ahead CSDA: the radial dependence of the dose, the longitudinal fluctuations in deposition (energy loss straggling), and non-straight trajectories (geometrical straggling).

B.1 Energy loss fluctuation on small thicknesses

In TIARA's `Raytracer` module, energy losses are normally computed as the product of the LET by the traveled distance, provided the LET does not vary significantly. However, this estimate is only an average (expectation) value, and in fact, due to the discrete nature of δ-rays, the energy deposited in a thin film actually fluctuates around this mean. Detailed Monte Carlo particle transport simulations are necessary to study energy distribution over lengths comparable to the mean free path of an ion in-between ionizing collisions. In [169], such a study was carried out, and the deviation in charge deposition Q from track to track over a certain thickness t_{Si}, was shown to depend on the moment ratio of the δ-rays' energy δ. Indeed:

$$Q = \frac{q}{E_{\mathrm{eh}}} \cdot \sum_{i=1}^{N_{\mathrm{colli}}} \delta_i$$

where we use our conventional notations, and N_{colli} is the number of collisions happening over thickness t_{Si}. Thus, under independent, identically distributed energy losses δ_i,

$$\bar{Q^2} = \frac{q^2}{E_{\mathrm{eh}}^2} \cdot N_{\mathrm{colli}} \cdot \bar{\delta^2} = \frac{q^2}{E_{\mathrm{eh}}^2} \cdot \left(\frac{\rho_{\mathrm{Si}} \cdot LET \cdot t_{\mathrm{Si}}}{\bar{\delta}} \right) \cdot \bar{\delta^2}$$

where in condensed notations $\bar{X} = E(X)$ stands for the average (expectation value), and $\bar{X^2} = E(X^2) - E(X)^2$ denotes variance. The average energy lost by the ion, $\rho_{\mathrm{Si}} \cdot LET \cdot t_{\mathrm{Si}}$, corresponds to a CSDA estimate as computed by TIARA's `Raytracer`, and ultimately,

calculating the deviation in deposited charge:

$$\sqrt{\bar{Q^2}} = \sqrt{\mathrm{cst} \cdot \bar{Q} \cdot \frac{\bar{\delta^2}}{\bar{\delta}}} \tag{B.1}$$

only amounts to determining the ratio of the second to the first moment of the δ-rays' energy. This quantity was estimated by the authors of [169], and fitted to a power law in the ion energy per unit mass E_{ion}:

$$\frac{\bar{\delta^2}}{\bar{\delta}} = \alpha \cdot E_{\mathrm{ion}}^{\beta} \tag{B.2}$$

In a subsequent paper, the same authors proposed new α and β parameters [170]. In more recent work [171, 172], the validity of this approach is confirmed, and to account for energy loss straggling the authors mostly propose different fitting coefficients.

To the best of the author's knowledge, this is as far as *analytical* formulas go in the current literature to describe charge deposition fluctuations, and thus for lack of higher-order expansions, we implemented (B.1) and (B.2) in TIARA for our SOI simulations. In practice, the `Raytracer`'s initial CSDA calculations are altered *a posteriori* for active-silicon track segments, by rescaling their charge content by a common amount computed from (B.1) and (B.2) assuming the overall traversed thickness is $t_{\mathrm{Si}} = z_{\mathrm{ref}}/|\cos\theta|$, where θ is the ion tilt and z_{ref} is a uniform thickness representative of the technology, e.g. 10 nm in 28 nm FD-SOI. Indeed, if fluctuations were computed on a per-segment basis, depending on track discretization the computed fluctuations could be overly large, e.g. for segments lying near the corners of a box. Given the computed mean and deviation in deposited charge for this t_{Si}, the charge-rescaling factor is then found by drawing a random number from a gamma distribution of matching moments (i.e. a straightforward probability law for a positive variate of independent mean and variance).

Figure B.1a shows how this translates into a Probability Density Function (PDF) of LETs for ion species used at RADEF and for a thickness of 10 nm. The spread is proportional to the root of the average and at small LETs, the mean, the median, and the mode (maximum likelihood) can be quite far apart, which means that the concept of average LET is not necessarily the prime parameter to describe charge deposition. In Figure B.1b, we plot heavy-ion simulation results on the SPHD bit-cell discussed in Chapter 4, with and without these fluctuations. As can be seen, the impact can be significant in low injection, with SEUs appearing at a lower threshold LET when fluctuations are activated, due to above-average events. On the contrary, at high LET the predicted cross section does not change, because even below-average events largely exceed the bit-cell's critical charge. Briefly speaking, accounting for longitudinal ion track fluctuations allows to refine the simulation accuracy near LET_{th}—here, capturing cross sections with nanometric resolution—with unchanged high-LET behavior and σ_{sat}. Similar conclusions were reached in [109].

B.2 Radial dose profile

Looking at track structure on the transverse dimensions now, due to the finite range of δ-rays, energy is transferred within a certain radius from the core of the ion track. Once again, this means that the sole LET parameter is not enough to characterize charge deposition, since two ions with equal LET but different energy and mass can exhibit

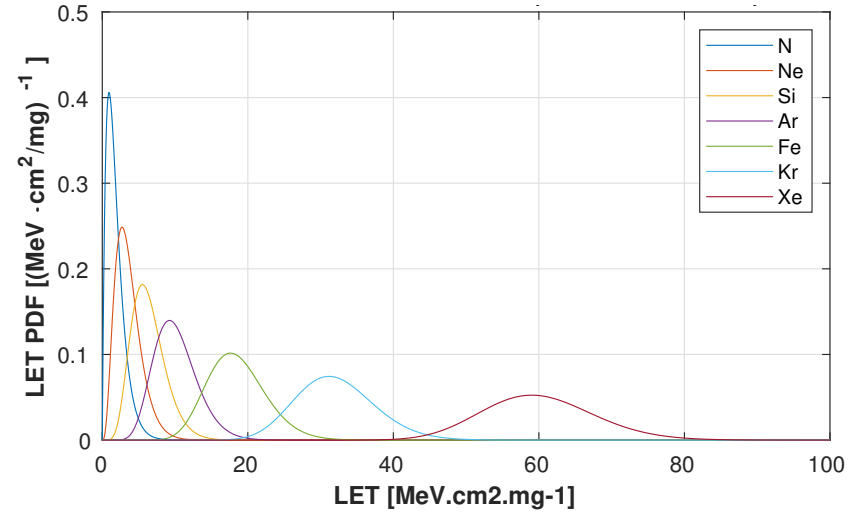

(a)

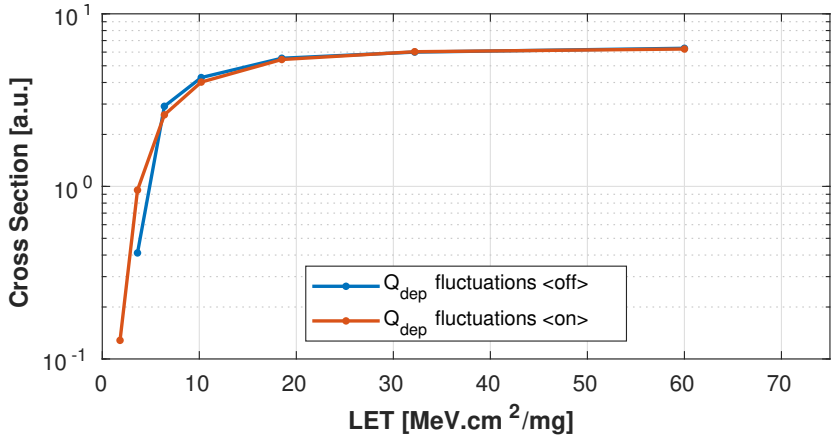

(b)

Figure B.1: Energy loss straggling in thin films: (a) LET fluctuations computed for a 10 nm thickness on RADEF 9.3 MeV/amu cocktail (average: 1.83, 3.63, 6.4, 10.2, 18.5, 32.2, 60 MeV·cm^2/mg) – (b) Effect on cross section predictions for a 28 nm FD-SOI SRAM.

different radial track profiles. The circuit cross section under these ions can then differ, especially for highly-integrated nodes and especially in SOI, where σ is in fact a reflection of a small sensitive area *convolved* with the ion's typical radius. Advanced Monte Carlo simulations are necessary to transport not only energetic ions, but also delta rays all the way down to thermalization as they locally ionize the target material—something which is described by the complex dielectric function theory. One of the most detailed works in that respect was carried out in [28], in which the authors were able to follow secondary electrons' trajectories down to a very low cut-off energy of 1.5 eV, i.e. where detailed semiconductor band structure also intervenes.

In terms of more behavioral modeling, the formulas published in [171] should be cited, showing two main regimes for the radial dose $D(r)$: most of the energy is deposited over small radii, e.g. more than half the dose is typically concentrated within the first few nanometers, and on longer scales $D(r)$ goes as $1/r^{2+\cdots}$, i.e. it decays a bit steeper than an inverse-square law. $D(r)$ then goes to true zero at the maximum δ-ray range, which can be as large as a few micrometers. This means that no *single* radius can be assigned to the ion track, and for instance in TCAD simulations accounting for a radially Gaussian

carrier generation term, either the track's core (say, 50% energy) or its full extent (100% energy) may be captured, but not both. Note that the magnitude of $D(r)$ depends on the LET (thus it roughly scales as Z^2/E), while the *shape* of $D(r)$ is mainly driven by the ion's energy over its mass; two ions with different charge but comparable energy per amu (or equivalently, velocity) will induce somewhat proportional ionizing tracks.

At present in TIARA, ray tracing calculations for charge deposition are done assuming infinitely-thin ion tracks. In Section 3.3.5, page 95, we mentioned that the generation term used for the CaTHODE 1D semiconductor simulation needs to smooth out the raw ray-traced outputs in order to avoid numerical divergences. Formally, $Q_s(x)$ a smooth distribution of charge in the 1D domain oriented on the x axis, is calculated by convolving the raw charge $Q(x)$ as deposited by the `Raytracer`, with a certain kernel $g(x)$:

$$Q_s(x_i) = \sum_j Q(x_j) \cdot g(x_i - x_j) \cdot \frac{x_{j+1} - x_{j-1}}{2} \tag{B.3}$$

This means that we can emulate the effect of a certain track profile, for instance at normal incidence if $g(x)$ is shaped as $\int_y D(r)\mathrm{d}y$, i.e. the track integrated over the perpendicular dimension (along transistor width). If the thin ray-traced track lies outside active silicon though, zero charge will still be deposited.

To avoid this, in future developments we will have to upgrade TIARA's representation of ion tracks so as to ray trace a bundle of lines instead of a single ray: with a sufficient number of parallel segments to carry the radial distribution of energy, we will be able to capture how the circuit cross section "bleeds" around the edges of active silicon in SOI. Mathematically, this means that the parallel lines must sample $D(r)$ adequately, and both deterministic or Monte Carlo approaches can be envisaged to distribute the lines over a certain radial grid or via importance sampling. A rather algorithmic difficulty may arise if the `Raytracer` is slowed down by having to intersect many segments against the circuit geometry: 3D structures in TIARA are now routinely comprised of several thousands of boxes, and for a fair representation of $D(r)$, dozens to hundreds of lines may be necessary. Given that the 3D structure is traversed in linear time for intersection checks, ray tracing operations could take up non-negligible time, if not become a bottleneck in the simulation. In that case, the collection of boxes will have to be arranged following a Bounding Volume Hierarchy (BVH), or a tree structure that accelerates intersection checks. Given the Manhattan geometries in TIARA, the most natural structure to implement is an Axis-Aligned Bounding Box (AABB) BVH, as il-

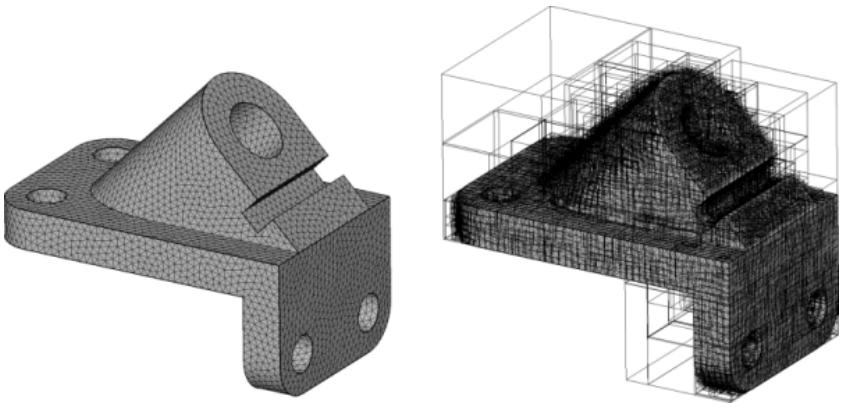

Figure B.2: Axis-Aligned Bounding Box (AABB) hierarchy for fast intersection detection [173].

lustrated in Figure B.2. Rather than checking a segment for intersections against all physical boxes (linear array), boxes can be (recursively, tree-wise) embedded in larger bounding boxes, and when the segment does not intersect a bounding box, all tests against the children boxes can be discarded at once. The complexity of intersection check then becomes essentially logarithmic (tree traversal) instead of linear, alleviating much of the computational burden.

B.3 Straggling of light particle trajectories

The final aspect of ion track structure omitted by straight-ahead CSDA transport, is the geometrical straggling of individual ion trajectories. As briefly discussed in Section 4.3.2, page 129, this can be relevant for light particles such as low-energy protons and also electrons, whose trajectories can be quite erratic near stopping. As depicted in Figure B.3, the angular distribution of protons close to the Bragg peak (i.e. 55 keV) is not trivial. With the CSDA ray tracing calculations, Proton Direct Ionization (PDI) is accounted for *on average*, but this yields a very abrupt upset threshold for PDI-induced upsets, associated with a large cross section. Experimental measurements as shown in Figure 4.24, page 128, on the other hand, demonstrate that low-energy proton upsets are quite rare, much more than upsets induced by low-LET heavy ions: only trajectories that are very *bent* lead to significant charge deposition in the sensitive volumes, which happens only for the small fraction of outliers in Figure B.3. For a practical modeling approach, TIARA will probably have to rely on a dictionary of trajectories: simple simulation via an "average ± deviation" scheme is not expected to yield correct results for PDI-induced upsets, since it will not truly capture the rare trajectories. In other words, the `Raytracer` calculations will have to be bypassed to load a trajectory—along with the associated energy losses—from a precomputed database. Note that for more detailed transport simulations taking advantage of TIARA's detailed 3D geometries and materials, interfacing with a code such as Geant4 may be more justified.

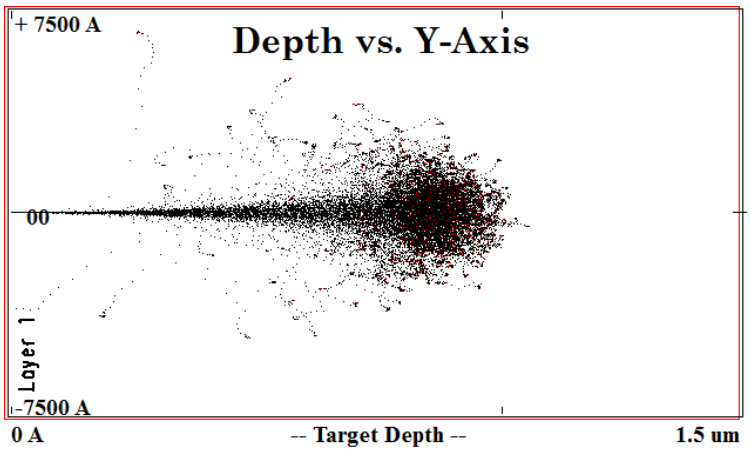

Figure B.3: One thousand 100 keV protons transported in silicon with SRIM.

Appendix C

Mathematical details of device simulation in CaTHODE

Our goal here is to shed more light onto the mathematical details of the discretization procedure that was described succinctly in Section 3.3.2, page 90. Suppose we want to compute the solution to our coupled PDE set on discrete grid points $\{x_i\}_{i=0..I-1}$ and discrete instants $\{t_j\}_{j=0..J-1}$. We will note $f(x_i, t_j) = f_{i,j}$ for short. We also scale our units to work on dimensionless quantities: all potentials are divided by the thermal voltage $V_T = k_B T/q$, all densities (carriers, doping, radiation...) are divided by the intrinsic density n_i, and to close the loop, all times are divided by an arbitrary, extrinsic, reference time. Using this unit-scaling paradigm, reference values for all other physical quantities can be derived as a combination of these primary scaling factors, and thus all equations can be nondimensionalized.

C.1 Spatial discretization

First, the box method [140] is used for spatial discretization; for all equations of the form $\nabla \cdot \mathcal{F} + S = 0$ i.e. (3.9), (3.10) and (3.11) on page 89, the idea is to use Ostrogradsky's theorem on a local control volume (the "box" depicted in Figure 3.6, page 91), yielding:

$$\sigma_{i+}\mathcal{F}_{i+} - \sigma_{i-}\mathcal{F}_{i-} + v_i S_i = 0 \tag{C.1}$$

In Poisson's equation for instance, this replaces the usual finite-difference formula for ΔV with a reweighted Laplacian operator based on the local mesh cross section, natively leading to steeper potential drops in areas where the silicon is thinner. More precisely, if we denote:

$$dx^2_{i-} = v_i(x_i - x_{i-1})/\sigma_{i-} \tag{C.2}$$

$$dx^2_{i+} = v_i(x_{i+1} - x_i)/\sigma_{i+} \tag{C.3}$$

then the Laplacian stencil is expressed as:

$$\Delta_i = \left[\frac{1}{dx^2_{i-}}, -\left(\frac{1}{dx^2_{i-}} + \frac{1}{dx^2_{i+}}\right), \frac{1}{dx^2_{i+}}\right] \tag{C.4}$$

Now the current equations (3.12) and (3.13) are more specific, and need to be discretized via the Scharfetter–Gummel scheme [141]. Briefly summarized, since the electron and hole densities have exponential dependence on the potential, when using naive finite-difference formulas the equations are made locally inconsistent unless the mesh is extremely dense. The Scharfetter–Gummel method prevents this issue by anticipating the values of n and p in-between grid points, yielding e.g. for the leftward current densities:

$$\mathcal{J}_{n,i-} = \frac{\mu_{n,i-}}{x_i - x_{i-1}} \left[n_i B(V_{n,i} - V_{n,i-1}) - n_{i-1} B(V_{n,i-1} - V_{n,i}) \right] \tag{C.5}$$

$$\mathcal{J}_{p,i-} = \frac{\mu_{p,i-}}{x_i - x_{i-1}} \left[p_{i-1} B(V_{p,i} - V_{n,i-1}) - p_i B(V_{p,i-1} - V_{p,i}) \right] \tag{C.6}$$

where the Bernoulli function:

$$B(x) = \frac{x}{\mathrm{e}^x - 1} \tag{C.7}$$

performs an optimal transition from a usual finite difference formula (which works well for diffusion-driven transport) to an upwind difference formula (which works well for drift-dominated problems) depending on carrier velocity. Note that $B(x)$ needs to be evaluated carefully to avoid overflows at large values of x and rounding errors due to the indeterminate form when $x \to 0$ (in that case we use a Taylor series expansion evaluated to machine precision).

C.2 Temporal discretization

As for temporal discretization, with the backward-Euler method all expressions of the form:

$$\partial g / \partial t + h = 0 \tag{C.8}$$

are replaced by:

$$(g_j - g_{j-1})/(t_j - t_{j-1}) + h_j = 0 \tag{C.9}$$

Then when going from $j-1$ to j, assembling all discretized equations (with (3.12) and (3.13) substituted into (3.10) and (3.11)) yields a nonlinear system of $3I$ equations for $3I$ unknowns—which we shall refrain from writing here in its entirety for the sake of compactness; if we denote our state vector X, then the system is of the form:

$$0_{3I} = f(V_{0,j}..V_{I-1,j}, n_{0,j}..n_{I-1,j}, p_{0,j}..p_{I-1,j}) = f(X_j) \tag{C.10}$$

The root to the equation residuals f is computed via the Newton–Raphson method, which performs successive refinements to the solution by linearizing the system. To go from "outer" iteration k to $k+1$ we compute an update candidate dX_j^{k+1} as the solution to the linearized system:

$$0_{3I} = J_f \cdot dX_j^{k+1} + f(X_j^k) \tag{C.11}$$

where J_f stands for the Jacobian matrix of f. Then to actually refresh the solution we compute:

$$X_j^{k+1} = X_j^k + \lambda^{k+1} dX_j^{k+1} \tag{C.12}$$

where λ is a damping parameter taken "as close as possible to 1", based on whether or not the update actually does improve the residuals—typically with geometrical damping for the "inner" iterations. Overall convergence is checked by ensuring that the norm of

the right-hand side falls below a given value. Optionally, we could also use a criterion on the norm of the solution update, but since the Newton method converges quadratically (doubling the number of exact digits at each step) when initialized close enough to the solution, this could cause to systematically perform one useless iteration just to let the solution stabilize when it is already good enough.

C.3 Equivalent conductance calculation

A last remaining mathematical aspect is the calculation of the equivalent conductance for our simulated device. It is defined as

$$G_{\text{eq}} = \partial\mathcal{I}/\partial U \tag{C.13}$$

where the applied voltage difference is:

$$U = V_l - V_r - (V_{0,\text{eq}} - V_{I-1,\text{eq}}) \tag{C.14}$$

and current $\mathcal{I}$ is obtained via the mesh average of the flux of the current densities in order to minimize numerical error:

$$\mathcal{I} = \frac{1}{I-1}\sum_{i=0}^{I-2}(\mathcal{J}_{n,i+} + \mathcal{J}_{p,i+}) \cdot \sigma_{i+} \tag{C.15}$$

(the displacement current $\mathcal{I}_{\text{d}}$ can safely be ignored: even when the voltage drops by 1 V in 1 ps, for our $L \approx 100$ nm, $\sigma \approx 10\times100$ nm^2 meshes it only reaches $\mathcal{I}_{\text{d}} \approx \varepsilon\sigma/L\cdot\partial U/\partial t \approx$ 1 µA). Intuitively, the derivative for G_{eq} should not have to be computed by evaluating the solution "at $U + dU$"—which could double the overall computation time—since the Jacobian matrix J_f already contains all the information about how the equation residuals vary around the solution. This is precisely the idea in [142]: we can write

$$G_{\text{eq}} = \frac{\partial\mathcal{I}}{\partial X}\cdot\frac{\partial X}{\partial U}$$

The first term is evaluated from symbolic differentiation, and to obtain the second term we make use of the fact that

$$f(X(U), U) = 0 \implies J_f \cdot \frac{\partial X}{\partial U} = -\frac{\partial f}{\partial U}$$

$\partial f/\partial U$ is a nearly-empty vector since U only appears at the boundary nodes, and because the Jacobian is already LU-factorized, the above system can be solved for $\partial X/\partial U$ at almost no computational overcost.

Appendix D

Personal publications

Peer-reviewed journals

Accurate Resolution of Time-Dependent and Circuit-Coupled Charge Transport Equations: 1D Case Applied to 28 nm FD-SOI Devices

Victor Malherbe, Gilles Gasiot, Thomas Thery, Jean-Luc Autran, Philippe Roche

IEEE Transactions on Nuclear Science, Jan. 2018 – Presented at NSREC 2017

On-Orbit Upset Rate Prediction at Advanced Technology Nodes: a 28 nm FD-SOI Case Study

Victor Malherbe, Gilles Gasiot, Dimitri Soussan, Jean-Luc Autran, Philippe Roche

IEEE Transactions on Nuclear Science, Jan. 2017 – Presented at NSREC 2016

Conference proceedings

Investigating the Single-Event-Transient Sensitivity of 65 nm Clock Trees with Heavy-Ion Irradiation and Monte-Carlo Simulation

Victor Malherbe, Gilles Gasiot, Sylvain Clerc, Fady Abouzeid, Jean-Luc Autran, Philippe Roche

IEEE International Reliability Physics Symposium, Apr. 2016

Alpha Soft Error Rate of FDSOI 28 nm SRAMs: Experimental Testing and Simulation Analysis

Victor Malherbe, Gilles Gasiot, Dimitri Soussan, Aurélien Patris, Jean-Luc Autran, Philippe Roche

IEEE International Reliability Physics Symposium, Apr. 2015

Muons and thermal neutrons SEU characterization of 28nm UTBB FD-SOI and Bulk eSRAMs

Gilles Gasiot, Dimitri Soussan, Jean-Luc Autran, Victor Malherbe, Philippe Roche

IEEE International Reliability Physics Symposium, Apr. 2015

Patents

Radiation-hardened CMOS logic device

Gilles Gasiot, Victor Malherbe, Sylvain Clerc

US patent #US9748955B1, Aug. 2017

Books

Charge Collection Physical Modeling for Soft Error Rate Computational Simulation

Jean-Luc Autran, Daniela Munteanu, Soilihi Moindjie, Tarek Saad Saoud, Victor Malherbe, Gilles Gasiot, Sylvain Clerc, Philippe Roche

in Digital Circuits, Modeling and Simulation in Engineering Sciences, Aug. 2016 – Noreen Sher Akbar and O. Anwar Beg, IntechOpen

Bibliography

[1] "Geant4 overview," http://www.geant4.org/, accessed 2018-10-22. [pp. vii and 12.]

[2] "Los Alamos National Laboratory: MCNP Home Page," http://mcnp.lanl.gov/, accessed 2018-10-22. [pp. vii and 12.]

[3] R. A. Weller, A. L. Sternberg, L. W. Massengill, R. D. Schrimpf, and D. M. Fleetwood, "Evaluating average and atypical response in radiation effects simulations," *IEEE Transactions on Nuclear Science*, vol. 50, no. 6, pp. 2265–2271, Dec 2003. [pp. vii, 36, and 75.]

[4] F. Wrobel and F. Saigné, "MC-ORACLE: A tool for predicting Soft Error Rate," *Computer Physics Communications*, vol. 182, no. 2, pp. 317 – 321, 2011. [pp. vii, 36, and 75.]

[5] S. Uznanski, "Monte-Carlo simulation and contribution to understanding of Single-Event-Upset (SEU) mechanisms in CMOS technologies down to 20nm technological node," Ph.D. dissertation, Université de Provence Aix-Marseille 1, 2011. [pp. viii, 46, 56, 70, and 116.]

[6] "The Stopping and Range of Ions in Matter Software," http://www.srim.org, accessed 2017-12-21. [pp. x and 13.]

[7] J.-M. Palau, G. Hubert, K. Coulie, B. Sagnes, M.-C. Calvet, and S. Fourtine, "Device simulation study of the SEU sensitivity of SRAMs to internal ion tracks generated by nuclear reactions," *IEEE Transactions on Nuclear Science*, vol. 48, no. 2, pp. 225–231, Apr 2001. [pp. x, 36, 45, 70, 75, 88, and 116.]

[8] "Eldo Platform," http://www.mentor.com/products/ic_nanometer_design/analog-mixed-signal-verification/eldo-platform, accessed 2018-03-27. [pp. x, 22, 51, 71, and 88.]

[9] "IBM Platform LSF," http://www.ibm.com/support/knowledgecenter/en/SSETD4/product_welcome_platform_lsf.html, accessed 2018-08-13. [pp. x and 51.]

[10] V. Malherbe, G. Gasiot, D. Soussan, A. Patris, J.-L. Autran, and P. Roche, "Alpha soft error rate of FDSOI 28 nm SRAMs: Experimental testing and simulation analysis," in *2015 IEEE Internationl Reliability Physics Symposium (IRPS)*, April 2015, pp. SE.11.1–SE.11.6. [pp. x, 80, 85, and 104.]

[11] P. Dodd, M. Shaneyfelt, K. Horn, D. Walsh, G. Hash, T. Hill, B. Draper, J. Schwank, F. Sexton, and P. Winokur, "SEU-sensitive volumes in bulk and SOI SRAMs from first-principles calculations and experiments," *IEEE Transactions on Nuclear Science*, vol. 48, no. 6, pp. 1893–1903, Dec 2001. [pp. x, 23, 35, 36, and 102.]

[12] J. R. Schwank, V. Ferlet-Cavrois, M. R. Shaneyfelt, P. Paillet, and P. E. Dodd, "Radiation effects in SOI technologies," *IEEE Transactions on Nuclear Science*, vol. 50, no. 3, pp. 522–538, June 2003. [pp. x and 80.]

[13] V. Malherbe, G. Gasiot, D. Soussan, J. L. Autran, and P. Roche, "On-Orbit Upset Rate Prediction at Advanced Technology Nodes: a 28 nm FD-SOI Case Study," *IEEE Transactions on Nuclear Science*, vol. 64, no. 1, pp. 449–456, Jan 2017. [pp. xi, 84, 99, and 121.]

[14] V. Malherbe, G. Gasiot, T. Thery, J. Autran, and P. Roche, "Accurate Resolution of Time-Dependent and Circuit-Coupled Charge Transport Equations: 1-D Case Applied to 28-nm FD-SOI Devices," *IEEE Transactions on Nuclear Science*, vol. 65, no. 1, pp. 331–338, Jan 2018. [pp. xii, 87, and 109.]

[15] G. Guennebaud, B. Jacob *et al.*, "Eigen v3," http://eigen.tuxfamily.org, 2010, accessed 2018-06-08. [pp. xii, 50, and 92.]

[16] V. Malherbe, G. Gasiot, S. Clerc, F. Abouzeid, J. Autran, and P. Roche, "Investigating the Single-Event-Transient Sensitivity of 65 nm Clock Trees with Heavy Ion Irradiation and Monte Carlo Simulation," in *2016 IEEE International Reliability Physics Symposium (IRPS)*, April 2016, pp. SE–3–1–SE–3–5. [pp. xiv and 109.]

[17] V. F. Hess, "Über Beobachtungen der durchdringenden Strahlung bei sieben Freiballonfahrten," *Physikalische Zeitschrift*, vol. 13, pp. 1084–1091, 1912. [p. 1.]

[18] "CREME-MC site," http://creme.isde.vanderbilt.edu/CREME-MC/, accessed 2018-10-22. [pp. 7 and 122.]

[19] "SPENVIS - Space Environment, Effects, and Education System," http://www.spenvis.oma.be/, accessed 2018-10-22. [pp. 7 and 122.]

[20] S. M. Seltzer, D. T. Bartlett, D. T. Burns, G. Dietze, H. Menzel, H. Paretzke, and A. Wambersie, "Fundamental Quantities and Units for Ionizing Radiation," *Journal of the International Commission on Radiation Units and Measurements*, vol. 11, Apr 2011. [pp. 8 and 31.]

[21] J. Barth, "Evolution of the Radiation Environment," in *2009 IEEE RADiation Effects on Components and Systems*, Sep 2009, short course presentation. [p. 9.]

[22] *Preliminary SEU analysis of the SAMPEX MIL-STD-1773 space-flight data*, vol. 1953, 1993. [p. 10.]

[23] K. Lilja, "Environment & Devices SER – Modeling Neutrons & Heavy Ion SER, from Planar CMOS to FinFETs," in *2016 IEEE Nuclear and Space Radiation Effects Conference*, July 2016, short course notebook, section III. [pp. 11 and 75.]

[24] "The official FLUKA site," http://www.fluka.org/, accessed 2018-10-22. [p. 12.]

[25] F. Wrobel, J.-M. Palau, M. C. Calvet, O. Bersillon, and H. Duarte, "Incidence of multi-particle events on soft error rates caused by n-Si nuclear reactions," *IEEE Transactions on Nuclear Science*, vol. 47, no. 6, pp. 2580–2585, Dec 2000. [p. 12.]

[26] "ENDF/B-VII.1 Evaluated Nuclear Data Library ," http://www.nndc.bnl.gov/endf/b7.1/, accessed 2018-10-22. [p. 12.]

[27] J. Ziegler, J. Biersack, and M. Ziegler, *SRIM: The Stopping and Range of Ions in Matter*, 2008. [p. 13.]

[28] M. Murat, A. Akkerman, and J. Barak, "Electron and Ion Tracks in Silicon: Spatial and Temporal Evolution," *IEEE Transactions on Nuclear Science*, vol. 55, no. 6, pp. 3046–3054, Dec 2008. [pp. 14 and 149.]

[29] R. C. Alig and S. Bloom, "Electron-Hole-Pair Creation Energies in Semiconductors," *Physical Review Letters*, vol. 35, pp. 1522–1525, Dec 1975. [p. 14.]

[30] M. Raine, "Étude de l'effet de l'énergie des ions lourds sur la sensibilité des composants électroniques," Ph.D. dissertation, Université Paris 11, 2011. [p. 15.]

[31] "Sentaurus Device," http://www.synopsys.com/silicon/tcad/device-simulation/sentaurus-device.html, accessed 2018-10-22. [pp. 17 and 80.]

[32] "SILVACO - Products - Device Simulation Framework," http://www.silvaco.com/products/tcad/device_simulation/device_simulation.html, accessed 2018-10-22. [p. 17.]

[33] F. B. McLean and T. R. Oldham, "Charge funneling in n-and p-type Si substrates," *IEEE Transactions on nuclear science*, vol. 29, no. 6, pp. 2017–2023, Dec 1982. [p. 18.]

[34] G. C. Messenger, "Collection of charge on junction nodes from ion tracks," *IEEE Transactions on Nuclear Science*, vol. 29, no. 6, pp. 2024–2031, Dec 1982. [p. 18.]

[35] "HSPICE – Synopsys," http://www.synopsys.com/verification/ams-verification/circuit-simulation/hspice.html, accessed 2018-03-27. [pp. 22, 51, and 71.]

[36] J.-L. Autran and D. Munteanu, *Soft Errors: From Particles to Circuits*. CRC Press, 2015. [pp. 22, 33, and 35.]

[37] "Latch-up – Wikipedia," http://en.wikipedia.org/wiki/Latch-up, accessed 2017-12-21. [p. 23.]

[38] J. T. Wallmark and S. M. Marcus, "Minimum Size and Maximum Packing Density of Nonredundant Semiconductor Devices," *Proceedings of the IRE*, vol. 50, no. 3, pp. 286–298, March 1962. [p. 23.]

[39] D. Binder, E. C. Smith, and A. B. Holman, "Satellite Anomalies from Galactic Cosmic Rays," *IEEE Transactions on Nuclear Science*, vol. 22, no. 6, pp. 2675–2680, Dec 1975. [p. 23.]

[40] P. Roche, J. M. Palau, G. Bruguier, C. Tavernier, R. Ecoffet, and J. Gasiot, "Determination of key parameters for SEU occurrence using 3-D full cell SRAM simulations," *IEEE Transactions on Nuclear Science*, vol. 46, no. 6, pp. 1354–1362, Dec 1999. [p. 24.]

[41] S. E. Diehl, J. E. Vinson, B. D. Shafer, and T. M. Mnich, "Considerations for Single Event Immune VLSI Logic," *IEEE Transactions on Nuclear Science*, vol. 30, no. 6, pp. 4501–4507, Dec 1983. [p. 25.]

[42] A. L. Friedman, B. Lawton, K. R. Hotelling, J. C. Pickel, V. H. Strahan, and K. Loree, "Single Event Upset in Combinatorial and Sequential Current Mode Logic," *IEEE Transactions on Nuclear Science*, vol. 32, no. 6, pp. 4216–4218, Dec 1985. [p. 25.]

[43] T. Karnik and P. Hazucha, "Characterization of soft errors caused by single event upsets in CMOS processes," *IEEE Transactions on Dependable and Secure Computing*, vol. 1, no. 2, pp. 128–143, April 2004. [p. 26.]

[44] J. L. Autran, S. Serre, D. Munteanu, S. Martinie, S. Semikh, S. Sauze, S. Uznanski, G. Gasiot, and P. Roche, "Real-time Soft-Error testing of 40nm SRAMs," in *2012 IEEE International Reliability Physics Symposium (IRPS)*, April 2012, pp. 3C.5.1–3C.5.9. [pp. 27 and 37.]

[45] "Pioneer Atom Splitter," *Radio-Craft*, vol. 18, no. 9, pp. 22–23, June 1947, accessed 2019-02-13. [Online]. Available: http://www.americanradiohistory.com/Archive-Radio-Craft/1940s/Radio-Craft-1947-Jun.pdf [p. 27.]

[46] A. Virtanen, "Radiation effects facility RADEF," in *2002 Proceedings of the Eighth IEEE International On-Line Testing Workshop*, 2002, p. 188. [pp. 27 and 100.]

[47] "RADiation Effects Facility – Department of Physics, University of Jyväskylä," http://www.jyu.fi/science/en/physics/research/infrastructures/accelerator-laboratory/radiation-effects-facility, accessed 2018-09-16. [pp. 27 and 100.]

[48] W. Hajdas, L. Adams, B. Nickson, and A. Zehnder, "The Proton Irradiation Facility at the Paul Scherrer Institute," *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms*, vol. 113, no. 1, pp. 54 – 58, 1996, accelerators in Applied Research and Technology. [pp. 28 and 103.]

[49] "Proton Irradiation Facility Home Page (PIF) – Paul Scherrer Institute," http://pif.web.psi.ch/, accessed 2018-09-18. [pp. 28 and 103.]

[50] "Single Event Effects Test Method and Guidelines – ESCC Basic Specification No. 25100," accessed 2018-09-16. [Online]. Available: https://escies.org/download/specdraftapppub?id=3095 [pp. 28 and 100.]

[51] V. Ferlet-Cavrois, P. Paillet, D. McMorrow, A. Torres, M. Gaillardin, J. S. Melinger, A. R. Knudson, A. B. Campbell, J. R. Schwank, G. Vizkelethy, M. R. Shaneyfelt, K. Hirose, O. Faynot, C. Jahan, and L. Tosti, "Direct measurement of transient pulses induced by laser and heavy ion irradiation in deca-nanometer devices," *IEEE Transactions on Nuclear Science*, vol. 52, no. 6, pp. 2104–2113, Dec 2005. [p. 29.]

[52] S. P. Buchner, F. Miller, V. Pouget, and D. P. McMorrow, "Pulsed-Laser Testing for Single-Event Effects Investigations," *IEEE Transactions on Nuclear Science*, vol. 60, no. 3, pp. 1852–1875, June 2013. [p. 30.]

[53] M. Nicolaidis and R. Perez, "Measuring the width of transient pulses induced by ionising radiation," in *2003 IEEE International Reliability Physics Symposium Proceedings*, March 2003, pp. 56–59. [p. 30.]

[54] B. Narasimham, V. Ramachandran, B. L. Bhuva, R. D. Schrimpf, A. F. Witulski, W. T. Holman, L. W. Massengill, J. D. Black, W. H. Robinson, and D. McMorrow, "On-chip characterization of single-event transient pulsewidths," *IEEE Transactions on Device and Materials Reliability*, vol. 6, no. 4, pp. 542–549, Dec 2006. [p. 30.]

[55] R. Harada, Y. Mitsuyama, M. Hashimoto, and T. Onoye, "Measurement circuits for acquiring SET pulsewidth distribution with sub-FO1-inverter-delay resolution," in *2010 11th International Symposium on Quality Electronic Design (ISQED)*, March 2010, pp. 839–844. [p. 30.]

[56] "Los Alamos Neutron Science Center," http://lansce.lanl.gov/, accessed 2018-10-27. [p. 30.]

[57] "Neutron Irradiation Facility — TRIUMF : Canada's particle accelerator center," http://www.triumf.ca/neutron-irradiation-facility, accessed 2018-10-27. [p. 30.]

[58] W. V. Roosbroeck, "Theory of the flow of electrons and holes in germanium and other semiconductors," *The Bell System Technical Journal*, vol. 29, no. 4, pp. 560–607, Oct 1950. [p. 34.]

[59] P. Roche, J. M. Palau, K. Belhaddad, G. Bruguier, R. Ecoffet, and J. Gasiot, "SEU response of an entire SRAM cell simulated as one contiguous three dimensional device domain," *IEEE Transactions on Nuclear Science*, vol. 45, no. 6, pp. 2534–2543, Dec 1998. [p. 34.]

[60] E. Petersen, J. Pickel, J. Adams, and E. Smith, "Rate prediction for single event effects-a critique," *IEEE Transactions on Nuclear Science*, vol. 39, no. 6, pp. 1577–1599, Dec 1992. [pp. 36, 76, and 122.]

[61] K. M. Warren, A. L. Sternberg, J. D. Black, R. A. Weller, R. A. Reed, M. H. Mendenhall, R. D. Schrimpf, and L. W. Massengill, "Heavy Ion Testing and Single Event Upset Rate Prediction Considerations for a DICE Flip-Flop," *IEEE Transactions on Nuclear Science*, vol. 56, no. 6, pp. 3130–3137, Dec 2009. [pp. 36 and 75.]

[62] V. Correas, F. Saigne, B. Sagnes, J. Boch, G. Gasiot, D. Giot, and P. Roche, "Simulation Tool for the Prediction of Heavy Ion Cross Section of Innovative 130-nm SRAMs," *IEEE Transactions on Nuclear Science*, vol. 55, no. 4, pp. 2036–2041, Aug 2008. [p. 36.]

[63] G. Hubert, S. Duzellier, C. Inguimbert, C. Boatella-Polo, F. Bezerra, and R. Ecoffet, "Operational SER Calculations on the SAC-C Orbit Using the Multi-Scales Single Event Phenomena Predictive Platform (MUSCASEP^3)," *IEEE Transactions on Nuclear Science*, vol. 56, no. 6, pp. 3032–3042, Dec 2009. [pp. 37 and 75.]

[64] L. Artola, G. Hubert, K. M. Warren, M. Gaillardin, R. D. Schrimpf, R. A. Reed, R. A. Weller, J. R. Ahlbin, P. Paillet, M. Raine, S. Girard, S. Duzellier, L. W. Massengill, and F. Bezerra, "SEU Prediction From SET Modeling Using Multi-Node Collection in Bulk Transistors and SRAMs Down to the 65 nm Technology Node," *IEEE Transactions on Nuclear Science*, vol. 58, no. 3, pp. 1338–1346, June 2011. [pp. 37 and 75.]

[65] R. A. Reed, R. A. Weller, A. Akkerman, J. Barak, W. Culpepper, S. Duzellier, C. Foster, M. Gaillardin, G. Hubert, T. Jordan, I. Jun, S. Koontz, F. Lei, P. McNulty, M. H. Mendenhall, M. Murat, P. Nieminen, P. O'Neill, M. Raine, B. Reddell, F. Saigné, G. Santin, L. Sihver, H. H. K. Tang, P. R. Truscott, and F. Wrobel, "Anthology of the Development of Radiation Transport Tools as Applied to Single Event Effects," *IEEE Transactions on Nuclear Science*, vol. 60, no. 3, pp. 1876–1911, June 2013. [pp. 37, 75, and 87.]

[66] S. Uznanski, G. Gasiot, P. Roche, C. Tavernier, and J.-L. Autran, "Single Event Upset and Multiple Cell Upset Modeling in Commercial Bulk 65-nm CMOS SRAMs and Flip-Flops," *IEEE Transactions on Nuclear Science*, vol. 57, no. 4, pp. 1876–1883, Aug 2010. [pp. 37 and 56.]

[67] P. Roche, G. Gasiot, J. Autran, D. Munteanu, R. Reed, and R. Weller, "Application of the TIARA Radiation Transport Tool to Single Event Effects Simulation," *IEEE Transactions on Nuclear Science*, vol. 61, no. 3, pp. 1498–1500, June 2014. [pp. 37 and 46.]

[68] J. S. Kauppila, A. L. Sternberg, M. L. Alles, A. M. Francis, J. Holmes, O. A. Amusan, and L. W. Massengill, "A Bias-Dependent Single-Event Compact Model Implemented Into BSIM4 and a 90 nm CMOS Process Design Kit," *IEEE Transactions on Nuclear Science*, vol. 56, no. 6, pp. 3152–3157, Dec 2009. [pp. 37, 75, 88, and 92.]

[69] N. Seifert, P. Slankard, M. Kirsch, B. Narasimham, V. Zia, C. Brookreson, A. Vo, S. Mitra, B. Gill, and J. Maiz, "Radiation-Induced Soft Error Rates of Advanced CMOS Bulk Devices," in *2006 IEEE International Reliability Physics Symposium Proceedings*, March 2006, pp. 217–225. [p. 38.]

[70] G. Gasiot, D. Giot, and P. Roche, "Multiple Cell Upsets as the Key Contribution to the Total SER of 65 nm CMOS SRAMs and Its Dependence on Well Engineering," *IEEE Transactions on Nuclear Science*, vol. 54, no. 6, pp. 2468–2473, Dec 2007. [pp. 38 and 39.]

[71] P. Roche, J.-L. Autran, G. Gasiot, and D. Munteanu, "Technology downscaling worsening radiation effects in bulk: SOI to the rescue," in *2013 IEEE International Electron Devices Meeting (IEDM)*, Dec 2013, pp. 31.1.1–31.1.4. [pp. 38, 41, 100, and 109.]

[72] T. Masson and R. Ferrant, "Memory insensitive to disturbances," US Patent US5 570 313A, 1996, accessed 2018-10-21. [Online]. Available: http://patents.google.com/patent/US5570313 [p. 38.]

[73] J. A. Maharrey, J. S. Kauppila, R. C. Harrington, P. Nsengiyumva, D. R. Ball, T. D. Haeffner, E. X. Zhang, B. L. Bhuva, W. T. Holman, and L. W. Massengill, "Impact of Single-Event Transient Duration and Electrical Delay at Reduced Supply Voltages on SET Mitigation Techniques," *IEEE Transactions on Nuclear Science*, vol. 65, no. 1, pp. 362–368, Jan 2018. [p. 39.]

[74] S. Buchner, M. Baze, D. Brown, D. McMorrow, and J. Melinger, "Comparison of error rates in combinational and sequential logic," *IEEE Transactions on Nuclear Science*, vol. 44, no. 6, pp. 2209–2216, Dec 1997. [pp. 39, 40, and 109.]

[75] G. Gasiot, D. Soussan, M. Glorieux, C. Bottoni, and P. Roche, "SER/SEL performances of SRAMs in UTBB FDSOI28 and comparisons with PDSOI and BULK counterparts," in *2014 IEEE International Reliability Physics Symposium*, June 2014, pp. SE.6.1–SE.6.5. [pp. 41, 100, and 109.]

[76] N. Seifert, B. Gill, S. Jahinuzzaman, J. Basile, V. Ambrose, Q. Shi, R. Allmon, and A. Bramnik, "Soft Error Susceptibilities of 22 nm Tri-Gate Devices," *IEEE Transactions on Nuclear Science*, vol. 59, no. 6, pp. 2666–2673, Dec 2012. [p. 41.]

[77] S. Lee, I. Kim, S. Ha, C. Yu, J. Noh, S. Pae, and J. Park, "Radiation-induced soft error rate analyses for 14 nm FinFET SRAM devices," in *2015 IEEE International Reliability Physics Symposium*, April 2015, pp. 4B.1.1–4B.1.4. [p. 41.]

[78] M. Swartwout, "Evolution of Satellite Mission Success," in *2017 IEEE Nuclear and Space Radiation Effects Conference*, July 2017, short course notebook, section III. [pp. 41 and 42.]

[79] P. Roche, "New Space and Automotive Industrial Radiation Paradigms," in *2018 IEEE RADiations Effects on Components and Systems*, Sep 2018, topical day keynote. [pp. 42, 44, and 109.]

[80] C. Henry, "FCC gets five new applications for non-geostationary satellite constellations," accessed 2018-11-01. [Online]. Available: http://spacenews.com/fcc-gets-five-new-applications-for-non-geostationary-satellite-constellations/ [p. 42.]

[81] P. Roche, G. Gasiot, S. Clerc, J.-M. Daveau, C. Bottoni, M. Glorieux, V. Huard, V. Dugoujon, F. Malou, and L. Hili, "A 65nm CMOS Platform for Space Applications: Qualification Test Results on Rad-Hard Microprocessors," in *2013 IEEE Nuclear and Space Radiation Effects Conference*, July 2013, poster presentation. [pp. 42 and 113.]

[82] "C65SPACE - Rad hard 65nm CMOS technology platform for space applications - STMicroelectronics," http://www.st.com/en/space-products/c65space.html, accessed 2018-10-22. [pp. 42 and 113.]

[83] M. Bruel, "Process for the production of thin semiconductor material films," US Patent US5 374 564A, 1991, accessed 2018-11-02. [Online]. Available: http://patents.google.com/patent/US5374564A/en [p. 43.]

[84] N. Planes, O. Weber, V. Barral, S. Haendler, D. Noblet, D. Croain, M. Bocat, P. Sassoulas, X. Federspiel, A. Cros, A. Bajolet, E. Richard, B. Dumont, P. Perreau, D. Petit, D. Golanski, C. Fenouillet-Beranger, N. Guillot, M. Rafik, V. Huard, S. Puget, X. Montagner, M.-A. Jaud, O. Rozeau, O. Saxod, F. Wacquant, F. Monsieur, D. Barge, L. Pinzelli, M. Mellier, F. Boeuf, F. Arnaud, and M. Haond, "28nm FDSOI technology platform for high-speed low-voltage digital applications," in *2012 Symposium on VLSI Technology (VLSIT)*, June 2012, pp. 133–134. [p. 43.]

[85] G. Gasiot, "Etude de la sensibilité de technologies CMOS/BULK et CMOS/SOI partiellement désertée très largement sub-microniques dans l'environnement radiatif terrestre," Ph.D. dissertation, Université Bordeaux 1, 2004. [p. 45.]

[86] V. Correas, "Prédiction de la section efficace des aléas logiques (SEU) et du taux d'aléas multiples (MCU) de technologies CMOS décananométriques sous irradiation aux ions lourds," Ph.D. dissertation, Université Montpellier 2, 2008. [p. 45.]

[87] B. D. Sierawski, B. Bhuva, R. Reed, N. Tam, B. Narasimham, K. Ishida, A. Hillier, M. Trinczek, E. Blackmore, S. Wen, and R. Wong, "Bias dependence of muon-induced single event upsets in 28 nm static random access memories," in *2014 IEEE International Reliability Physics Symposium*, June 2014, pp. 2B.2.1–2B.2.5. [p. 46.]

[88] J. M. Trippe, R. A. Reed, R. A. Austin, B. D. Sierawski, R. A. Weller, E. D. Funkhouser, M. P. King, B. Narasimham, B. Bartz, R. Baumann, J. Labello, J. Nichols, R. D. Schrimpf, and S. L. Weeden-Wright, "Electron-Induced Single Event Upsets in 28 nm and 45 nm Bulk SRAMs," *IEEE Transactions on Nuclear Science*, vol. 62, no. 6, pp. 2709–2716, Dec 2015. [p. 46.]

[89] G. Gasiot, D. Soussan, J. Autran, V. Malherbe, and P. Roche, "Muons and thermal neutrons SEU characterization of 28nm UTBB FD-SOI and Bulk eSRAMs," in *2015 IEEE International Reliability Physics Symposium*, April 2015, pp. 2C.2.1–2C.2.5. [p. 46.]

[90] N. Seifert, S. Jahinuzzaman, J. Velamala, and N. Patel, "Susceptibility of planar and 3D tri-gate technologies to muon-induced single event upsets," in *2015 IEEE International Reliability Physics Symposium*, April 2015, pp. 2C.1.1–2C.1.6. [p. 46.]

[91] "Boost C++ Libraries," http://www.boost.org, accessed 2018-09-15. [p. 50.]

[92] "Wavefront OBJ File Format Summary," http://www.fileformat.info/format/wavefrontobj/egff.htm, accessed 2018-09-15. [p. 51.]

[93] "Home of the Blender project – Free and Open 3D Creation Software," http://www.blender.org, accessed 2018-09-15. [p. 51.]

[94] "Graph isomorphism – Wikipedia," http://en.wikipedia.org/wiki/Graph_isomorphism, accessed 2018-04-23. [p. 54.]

[95] S. Uznanski, G. Gasiot, P. Roche, J.-L. Autran, and L. Dugoujon, "Heavy Ion Characterization and Monte Carlo Simulation on 32 nm CMOS Bulk Technology," *IEEE Transactions on Nuclear Science*, vol. 58, no. 6, pp. 2652–2657, Dec 2011. [p. 56.]

[96] M. Matsumoto and T. Nishimura, "Mersenne Twister: a 623-dimensionally equidistributed uniform pseudo-random number generator," *ACM Transactions on Modeling and Computer Simulation*, vol. 8, pp. 3–30, Jan 1998. [pp. 56 and 139.]

[97] "Mersenne Twister Home Page," http://www.math.sci.hiroshima-u.ac.jp/~m-mat/MT/emt.html, accessed 2018-08-17. [pp. 56 and 139.]

[98] J. Detwiler, R. Henning, R. Johnson, and M. Marino, "A Generic Surface Sampler for Monte Carlo Simulations," *IEEE Transactions on Nuclear Science*, vol. 55, pp. 2329 – 2333, Sept 2008. [p. 60.]

[99] J. Rovira, P. Wonka, F. Castro, and M. Sbert, "Point sampling with uniformly distributed lines," in *Proceedings Eurographics/IEEE VGTC Symposium Point-Based Graphics*, June 2005, pp. 109–118. [p. 60.]

[100] S. Serre, S. Semikh, S. Uznanski, J. L. Autran, D. Munteanu, G. Gasiot, and P. Roche, "Geant4 Analysis of n-Si Nuclear Reactions From Different Sources of Neutrons and Its Implication on Soft-Error Rate," *IEEE Transactions on Nuclear Science*, vol. 59, no. 4, pp. 714–722, Aug 2012. [p. 64.]

[101] K. M. Warren, "Monte Carlo Based Single-Event Effect and Soft-Error Rate Prediction Methods," in *2012 IEEE Nuclear and Space Radiation Effects Conference*, July 2012, short course notebook, section IV. [p. 66.]

[102] G. Gasiot, V. Ferlet-Cavrois, J. Baggio, P. Roche, P. Flatresse, A. Guyot, P. Morel, O. Bersillon, and J. du Port de Pontcharra, "SEU sensitivity of bulk and SOI technologies to 14-MeV neutrons," *IEEE Transactions on Nuclear Science*, vol. 49, no. 6, pp. 3032–3037, Dec 2002. [p. 66.]

[103] M. Glorieux, J. Autran, D. Munteanu, S. Clerc, G. Gasiot, and P. Roche, "Random-Walk Drift-Diffusion Charge-Collection Model for Reverse-Biased Junctions Embedded in Circuits," *IEEE Transactions on Nuclear Science*, vol. 61, no. 6, pp. 3527–3534, Dec 2014. [pp. 70, 71, and 92.]

[104] "Verilog-AMS Language Reference Manual," accessed 2018-09-15. [Online]. Available: http://www.accellera.org/images/downloads/standards/v-ams/VAMS-LRM-2-4.pdf [p. 72.]

[105] H. H. K. Tang, "SEMM-2: A new generation of single-event-effect modeling tools," *IBM Journal of Research and Development*, vol. 52, no. 3, pp. 233–244, May 2008. [p. 75.]

[106] H. H. K. Tang and E. H. Cannon, "SEMM-2: a modeling system for single event upset analysis," *IEEE Transactions on Nuclear Science*, vol. 51, no. 6, pp. 3342–3348, Dec 2004. [p. 75.]

[107] K. Foley, N. Seifert, J. B. Velamala, W. G. Bennett, and S. Gupta, "IRT: A modeling system for single event upset analysis that captures charge sharing effects," in *2014 IEEE International Reliability Physics Symposium*, June 2014, pp. 5F.1.1–5F.1.9. [p. 75.]

[108] S. Abe, Y. Watanabe, N. Shibano, N. Sano, H. Furuta, M. Tsutsui, T. Uemura, and T. Arakawa, "Multi-Scale Monte Carlo Simulation of Soft Errors Using PHITS-HyENEXSS Code System," *IEEE Transactions on Nuclear Science*, vol. 59, no. 4, pp. 965–970, Aug 2012. [p. 75.]

[109] M. Raine, G. Hubert, P. Paillet, M. Gaillardin, and A. Bournel, "Implementing Realistic Heavy Ion Tracks in a SEE Prediction Tool: Comparison Between Different Approaches," *IEEE Transactions on Nuclear Science*, vol. 59, no. 4, pp. 950–957, Aug 2012. [pp. 75 and 148.]

[110] K. Lilja, M. Bounasser, and T. Assis, "Single Event Simulation and Error Rate Prediction for Space Electronics in Advanced Semiconductor Technologies," in *6th International Workshop on Analogue and Mixed-Signal Integrated Circuits for Space Applications (AMICSA)*, June 2016. [p. 75.]

[111] H. Belhaddad and R. Perez, "Apparatus and method for the determination of SEU and SET disruptions in a circuit caused by ionizing particle strikes," US Patent US20 080 077 376A1, 2008. [p. 75.]

[112] S. DasGupta, A. F. Witulski, B. L. Bhuva, M. L. Alles, R. A. Reed, O. A. Amusan, J. R. Ahlbin, R. D. Schrimpf, and L. W. Massengill, "Effect of Well and Substrate Potential Modulation on Single Event Pulse Shape in Deep Submicron CMOS," *IEEE Transactions on Nuclear Science*, vol. 54, no. 6, pp. 2407–2412, Dec 2007. [pp. 76, 86, and 95.]

[113] V. Ferlet-Cavrois, P. Paillet, M. Gaillardin, D. Lambert, J. Baggio, J. R. Schwank, G. Vizkelethy, M. R. Shaneyfelt, K. Hirose, E. W. Blackmore, O. Faynot, C. Jahan,

and L. Tosti, "Statistical Analysis of the Charge Collected in SOI and Bulk Devices Under Heavy Ion and Proton Irradiation—Implications for Digital SETs," *IEEE Transactions on Nuclear Science*, vol. 53, no. 6, pp. 3242–3252, Dec 2006. [pp. 76, 86, and 95.]

[114] P. Dodd, "Device simulation of charge collection and single-event upset," *IEEE Transactions on Nuclear Science*, vol. 43, no. 2, pp. 561–575, Apr 1996. [pp. 80, 82, and 88.]

[115] E. P. V. Ploeg, C. T. Nguyen, S. S. Wong, and J. D. Plummer, "Parasitic bipolar gain in fully depleted n-channel SOI MOSFET's," *IEEE Transactions on Electron Devices*, vol. 41, no. 6, pp. 970–977, June 1994. [pp. 80 and 81.]

[116] S. Cristoloveanu, "Introduction to Silicon On Insulator materials and devices ," *Microelectronic Engineering*, vol. 39, no. 1–4, pp. 145 – 154, 1997. [p. 80.]

[117] V. Ferlet-Cavrois, P. Paillet, D. Mcmorrow, J. S. Melinger, A. B. Campbell, M. Gaillardin, O. Faynot, O. Thomas, G. Barna, and B. Giffard, "Analysis of the Transient Response of High Performance 50-nm Partially Depleted SOI Transistors Using a Laser Probing Technique," *IEEE Transactions on Nuclear Science*, vol. 53, no. 4, pp. 1825–1833, Aug 2006. [p. 80.]

[118] V. Ferlet-Cavrois, C. Marcandella, G. Giraud, G. Gasiot, I. Colladant, O. Musseau, C. Fenouillet, and J. du Port de Poncharra, "Characterization of the parasitic bipolar amplification in SOI technologies submitted to transient irradiation," *IEEE Transactions on Nuclear Science*, vol. 49, no. 3, pp. 1456–1461, June 2002. [p. 80.]

[119] V. Ferlet-Cavrois, G. Gasiot, C. Marcandella, C. D'Hose, O. Flament, O. Faynot, J. du Port de Pontcharra, and C. Raynaud, "Insights on the transient response of fully and partially depleted SOI technologies under heavy-ion and dose-rate irradiations," *IEEE Transactions on Nuclear Science*, vol. 49, no. 6, pp. 2948–2956, Dec 2002. [p. 80.]

[120] D. Munteanu, V. Ferlet-Cavrois, J. L. Autran, P. Paillet, J. Baggio, O. Faynot, C. Jahan, and L. Tosti, "Investigation of quantum effects in ultra-thin body single- and double-gate devices submitted to heavy ion irradiation," *IEEE Transactions on Nuclear Science*, vol. 53, no. 6, pp. 3363–3371, Dec 2006. [p. 80.]

[121] V. Correas, I. Nofal, J. Cerba, F. Monsieur, G. Gasiot, D. Alexandrescu, P. Roche, and R. Gonella, "Analysis and Modeling of the Charge Collection Mechanism in 28-nm FD-SOI," *IEEE Transactions on Nuclear Science*, vol. 65, no. 8, pp. 1894–1899, Aug 2018. [pp. 80 and 86.]

[122] J. del Alamo, S. Swirhun, and R. M. Swanson, "Simultaneous measurement of hole lifetime, hole mobility and bandgap narrowing in heavily doped n-type silicon," in *1985 International Electron Devices Meeting*, vol. 31, 1985, pp. 290–293. [pp. 81 and 89.]

[123] D. Klaassen, "A unified mobility model for device simulation—I. Model equations and concentration dependence ," *Solid-State Electronics*, vol. 35, no. 7, pp. 953 – 959, 1992. [p. 82.]

[124] J. S. Laird, T. Hirao, S. Onoda, and H. Itoh, "High-injection carrier dynamics generated by MeV heavy ions impacting high-speed photodetectors," *Journal of Applied Physics*, vol. 98, no. 1, pp. 013 530–013 530–14, July 2005. [p. 82.]

[125] L. Edmonds, "Charge collection from ion tracks in simple EPI diodes," *IEEE Transactions on Nuclear Science*, vol. 44, no. 3, pp. 1448–1463, June 1997. [p. 82.]

[126] R. V. Overstraeten and H. D. Man, "Measurement of the ionization rates in diffused silicon p-n junctions," *Solid-State Electronics*, vol. 13, no. 5, pp. 583 – 608, 1970. [pp. 82 and 90.]

[127] O. Musseau and V. Ferlet-Cavrois, "Silicon-on-Insulator Technologies: Radiation Effects," in *2001 IEEE Nuclear and Space Radiation Effects Conference*, July 2001, short course notebook, section III. [p. 84.]

[128] K. Castellani-Coulie, D. Munteanu, V. Ferlet-Cavrois, and J. Autran, "Simulation analysis of the bipolar amplification in fully-depleted SOI technologies under heavy-ion irradiations," *IEEE Transactions on Nuclear Science*, vol. 52, no. 5, pp. 1474–1479, October 2005. [pp. 84 and 124.]

[129] M. L. Alles, "SPICE analysis of the SEU sensitivity of a fully depleted SOI CMOS SRAM cell," *IEEE Transactions on Nuclear Science*, vol. 41, no. 6, pp. 2093–2097, Dec 1994. [p. 87.]

[130] L. D. Edmonds, "A Theoretical Analysis of Steady-State Charge Collection in Simple Diodes Under High-Injection Conditions," *IEEE Transactions on Nuclear Science*, vol. 57, no. 2, pp. 818–830, Apr 2010. [p. 87.]

[131] D. Fulkerson and E. Vogt, "Prediction of SOI single-event effects using a simple physics-based SPICE model," *IEEE Transactions on Nuclear Science*, vol. 52, no. 6, pp. 2168–2174, Dec 2005. [p. 88.]

[132] J. S. Kauppila, L. W. Massengill, D. R. Ball, M. L. Alles, R. D. Schrimpf, T. D. Loveless, J. A. Maharrey, R. C. Quinn, and J. D. Rowe, "Geometry-Aware Single-Event Enabled Compact Models for Sub-50 nm Partially Depleted Silicon-on-Insulator Technologies," *IEEE Transactions on Nuclear Science*, vol. 62, no. 4, pp. 1589–1598, Aug 2015. [p. 88.]

[133] W. Bludau, A. Onton, and W. Heinke, "Temperature dependence of the band gap of silicon," *Journal of Applied Physics*, vol. 45, no. 4, pp. 1846–1848, 1974. [p. 89.]

[134] M. A. Green, "Intrinsic concentration, effective densities of states, and effective mass in silicon," *Journal of Applied Physics*, vol. 67, no. 6, pp. 2944–2954, 1990. [p. 89.]

[135] J. E. Lang, F. L. Madarasz, and P. M. Hemenger, "Temperature dependent density of states effective mass in nonparabolic p-type silicon," *Journal of Applied Physics*, vol. 54, no. 6, pp. 3612–3612, 1983. [p. 89.]

[136] N. D. Arora, J. R. Hauser, and D. J. Roulston, "Electron and hole mobilities in silicon as a function of concentration and temperature," *IEEE Transactions on Electron Devices*, vol. 29, no. 2, pp. 292–295, Feb 1982. [p. 90.]

[137] C. Canali, G. Majni, R. Minder, and G. Ottaviani, "Electron and hole drift velocity measurements in silicon and their empirical relation to electric field and temperature," *IEEE Transactions on Electron Devices*, vol. 22, no. 11, pp. 1045–1047, Nov 1975. [p. 90.]

[138] S. C. Choo, "Theory of a forward-biased diffused-junction P-L-N rectifier – Part I: Exact numerical solutions," *IEEE Transactions on Electron Devices*, vol. 19, no. 8, pp. 954–966, Aug 1972. [p. 90.]

[139] N. H. Fletcher, "The High Current Limit for Semiconductor Junction Devices," *Proceedings of the IRE*, vol. 45, no. 6, pp. 862–872, June 1957. [p. 90.]

[140] R. E. Bank, D. J. Rose, and W. Fichtner, "Numerical methods for semiconductor device simulation," *IEEE Transactions on Electron Devices*, vol. 30, no. 9, pp. 1031–1041, 1983. [pp. 90 and 153.]

[141] D. L. Scharfetter and H. K. Gummel, "Large-signal analysis of a silicon read diode oscillator," *IEEE Transactions on Electron Devices*, vol. 16, no. 1, pp. 64–77, 1969. [pp. 90 and 154.]

[142] K. Mayaram and D. O. Pederson, "Coupling algorithms for mixed-level circuit and device simulation," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 11, no. 8, pp. 1003–1012, Aug 1992. [pp. 91 and 155.]

[143] "Heavy Ion Facility (HIF) – Université Catholique de Louvain," http://uclouvain.be/en/research-institutes/irmp/crc/heavy-ion-facility-hif.html, accessed 2018-09-18. [p. 102.]

[144] H. Zhang, H. Jiang, B. L. Bhuva, J. S. Kauppila, W. T. Holman, and L. W. Massengill, "Frequency Dependence of Heavy-Ion-Induced Single-Event Responses of Flip-Flops in a 16-nm Bulk FinFET Technology," *IEEE Transactions on Nuclear Science*, vol. 65, no. 1, pp. 413–417, Jan 2018. [p. 109.]

[145] C. Bottoni, M. Glorieux, J. M. Daveau, G. Gasiot, F. Abouzeid, S. Clerc, L. Naviner, and P. Roche, "Heavy ions test result on a 65nm Sparc-V8 radiation-hard microprocessor," in *2014 IEEE International Reliability Physics Symposium*, June 2014, pp. 5F.5.1–5F.5.6. [p. 109.]

[146] C. Bottoni, B. Coeffic, J. Daveau, L. Naviner, and P. Roche, "Partial triplication of a SPARC-V8 microprocessor using fault injection," in *2015 IEEE 6th Latin American Symposium on Circuits Systems (LASCAS)*, Feb 2015, pp. 1–4. [p. 109.]

[147] C. Bottoni, B. Coeffic, J.-M. Daveau, G. Gasiot, L. Naviner, and P. Roche, "A Layout-Aware Approach to Fault Injection for Improving Failure Mode Prediction," in *Silicon Error in Logic – System Effects*, Apr. 2015. [p. 109.]

[148] V. Ferlet-Cavrois, P. Paillet, D. McMorrow, N. Fel, J. Baggio, S. Girard, O. Duhamel, J. S. Melinger, M. Gaillardin, J. R. Schwank, P. E. Dodd, M. R. Shaneyfelt, and J. A. Felix, "New Insights Into Single Event Transient Propagation in Chains of Inverters—Evidence for Propagation-Induced Pulse Broadening," *IEEE Transactions on Nuclear Science*, vol. 54, no. 6, pp. 2338–2346, Dec 2007. [p. 109.]

[149] J. A. Maharrey, R. C. Quinn, T. D. Loveless, J. S. Kauppila, S. Jagannathan, N. M. Atkinson, N. J. Gaspard, E. X. Zhang, M. L. Alles, B. L. Bhuva, W. T. Holman, and L. W. Massengill, "Effect of Device Variants in 32 nm and 45 nm SOI on SET Pulse Distributions," *IEEE Transactions on Nuclear Science*, vol. 60, no. 6, pp. 4399–4404, Dec 2013. [p. 109.]

[150] O. A. Amusan, L. W. Massengill, B. L. Bhuva, S. DasGupta, A. F. Witulski, and J. R. Ahlbin, "Design Techniques to Reduce SET Pulse Widths in Deep-Submicron Combinational Logic," *IEEE Transactions on Nuclear Science*, vol. 54, no. 6, pp. 2060–2064, Dec 2007. [p. 109.]

[151] E. H. Cannon, A. Kleinosowski, K. P. Muller, T. H. Ning, P. J. Oldiges, L. J. Sigal, J. D. Warnock, and D. Wendel, "Apparatus and method for hardening latches in SOI CMOS devices," US Patent US7 888 959B2, 2007, accessed 2018-10-21. [Online]. Available: http://patents.google.com/patent/US7888959 [p. 110.]

[152] J. S. Kauppila, T. D. Loveless, R. C. Quinn, J. A. Maharrey, M. L. Alles, M. W. McCurdy, R. A. Reed, B. L. Bhuva, L. W. Massengill, and K. Lilja, "Utilizing device stacking for area efficient hardened SOI flip-flop designs," in *2014 IEEE International Reliability Physics Symposium*, June 2014, pp. SE.4.1–SE.4.7. [p. 110.]

[153] H. . Wang, J. S. Kauppila, K. Lilja, M. Bounasser, L. Chen, M. Newton, Y. . Li, R. Liu, B. L. Bhuva, S. . Wen, R. Wong, R. Fung, S. Baeg, and L. W. Massengill, "Evaluation of SEU Performance of 28-nm FDSOI Flip-Flop Designs," *IEEE Transactions on Nuclear Science*, vol. 64, no. 1, pp. 367–373, Jan 2017. [p. 110.]

[154] N. Seifert, P. Shipley, M. Pant, V. Ambrose, and B. Gill, "Radiation-induced clock jitter and race," in *IEEE International Reliability Physics Symposium (IRPS), 43rd Annual Proceedings*, April 2005, pp. 215–222. [pp. 113, 117, and 119.]

[155] L. Wissel, D. Heidel, M. Gordon, K. Rodbell, K. Stawiasz, and E. Cannon, "Flip-Flop Upsets From Single-Event-Transients in 65 nm Clock Circuits," *IEEE Transactions on Nuclear Science (TNS)*, vol. 56, no. 6, pp. 3145–3151, Dec 2009. [p. 113.]

[156] R. Garg and S. Khatri, "A novel, highly SEU tolerant digital circuit design approach," in *IEEE International Conference on Computer Design (ICCD)*, Oct 2008, pp. 14–20. [p. 113.]

[157] R. Dash, R. Garg, S. Khatri, and G. Choi, "SEU hardened clock regeneration circuits," in *International Symposium on Quality of Electronic Design (ISQED)*, March 2009, pp. 806–813. [p. 113.]

[158] R. Chipana and F. Lima Kastensmidt, "SET Susceptibility Analysis of Clock Tree and Clock Mesh Topologies," in *IEEE Computer Society Annual Symposium on VLSI (ISVLSI)*, July 2014, pp. 559–564. [p. 113.]

[159] J. M. Mogollon, F. R. Palomo, M. A. Aguirre, J. Napoles, H. Guzman-Miranda, and E. Garcia-Sanchez, "TCAD Simulations on CMOS Propagation Induced Pulse Broadening Effect: Dependence Analysis on the Threshold Voltage," *IEEE Transactions on Nuclear Science*, vol. 57, no. 4, pp. 1908–1914, Aug 2010. [p. 114.]

[160] M. Glorieux, A. Evans, V. Ferlet-Cavrois, C. Boatella-Polo, D. Alexandrescu, S. Clerc, G. Gasiot, and P. Roche, "Detailed SET Measurement and Characterization of a 65 nm Bulk Technology," *IEEE Transactions on Nuclear Science*, vol. 64, no. 1, pp. 81–88, Jan 2017. [p. 114.]

[161] M. Glorieux, A. Evans, D. Alexandrescu, C. Boatella-Polo, K. Sanchez, and V. Ferlet-Cavrois, "DAMSEL—Dynamic and Applicative Measurement of Single Events in Logic," *IEEE Transactions on Nuclear Science*, vol. 65, no. 1, pp. 354–361, Jan 2018. [p. 114.]

[162] A. Mallajosyula and P. Zarkesh-Ha, "A Robust Single Event Upset Hardened Clock Distribution Network," in *2008 IEEE International Integrated Reliability Workshop Final Report*, Oct 2008, pp. 121–124. [p. 121.]

[163] E. L. Petersen, P. Shapiro, J. H. Adams, and E. A. Burke, "Calculation of Cosmic-Ray Induced Soft Upsets and Scaling in VLSI Devices," *IEEE Transactions on Nuclear Science*, vol. 29, no. 6, pp. 2055–2063, Dec 1982. [p. 124.]

[164] R. A. Reed, R. A. Weller, M. H. Mendenhall, D. M. Fleetwood, K. M. Warren, B. D. Sierawski, M. P. King, R. D. Schrimpf, and E. C. Auden, "Physical Processes and Applications of the Monte Carlo Radiative Energy Deposition (MRED) Code," *IEEE Transactions on Nuclear Science*, vol. 62, no. 4, pp. 1441–1461, Aug 2015. [p. 124.]

[165] B. D. Sierawski, J. A. Pellish, R. A. Reed, R. D. Schrimpf, K. M. Warren, R. A. Weller, M. H. Mendenhall, J. D. Black, A. D. Tipton, M. A. Xapsos, R. C. Baumann, X. Deng, M. J. Campola, M. R. Friendlich, H. S. Kim, A. M. Phan, and C. M. Seidleck, "Impact of Low-Energy Proton Induced Upsets on Test Methods and Rate Predictions," *IEEE Transactions on Nuclear Science*, vol. 56, no. 6, pp. 3085–3092, Dec 2009. [p. 127.]

[166] N. A. Dodds, J. R. Schwank, M. R. Shaneyfelt, P. E. Dodd, B. L. Doyle, M. Trinczek, E. W. Blackmore, K. P. Rodbell, M. S. Gordon, R. A. Reed, J. A. Pellish, K. A. LaBel, P. W. Marshall, S. E. Swanson, G. Vizkelethy, S. V. Deusen, F. W. Sexton, and M. J. Martinez, "Hardness Assurance for Proton Direct Ionization-Induced SEEs Using a High-Energy Proton Beam," *IEEE Transactions on Nuclear Science*, vol. 61, no. 6, pp. 2904–2914, Dec 2014. [p. 127.]

[167] G. Gasiot, V. Malherbe, and S. Clerc, "Radiation-hardened CMOS logic device," US Patent US9 748 955B1, 2017, accessed 2018-10-21. [Online]. Available: http://patents.google.com/patent/US9748955 [p. 131.]

[168] "Global Illumination Compendium – The Concise Guide to Global Illumination Algorithms," accessed 2018-06-22. [Online]. Available: http://people.cs.kuleuven.be/~philip.dutre/GI/ [pp. 141 and 145.]

[169] M. A. Xapsos, "Applicability of LET to single events in microelectronic structures," *IEEE Transactions on Nuclear Science*, vol. 39, no. 6, pp. 1613–1621, Dec 1992. [pp. 147 and 148.]

[170] M. A. Xapsos, T. R. Weatherford, and R. Shapiro, "The shape of heavy ion upset cross section curves," *IEEE Transactions on Nuclear Science*, vol. 40, no. 6, pp. 1812–1819, Dec 1993. [p. 148.]

[171] A. Akkerman and J. Barak, "Ion-track structure and its effects in small size volumes of silicon," *IEEE Transactions on Nuclear Science*, vol. 49, no. 6, pp. 3022–3031, Dec 2002. [pp. 148 and 149.]

[172] ——, "Correction to 'Ion-track structure and its effects in small size volumes of silicon'," *IEEE Transactions on Nuclear Science*, vol. 50, no. 3, pp. 741–741, June 2003. [p. 148.]

[173] P. Alliez, S. Tayeb, and C. Wormser, "3D Fast Intersection and Distance Computation," in *CGAL User and Reference Manual*, 4th ed. CGAL

Editorial Board, 2018, accessed 2019-02-13. [Online]. Available: https://doc.cgal.org/4.13/Manual/packages.html#PkgAABB_treeSummary [p. 150.]