Migrating From Flat Boosts to Frame Trees: A Hierarchical Lorentz Approach for High-Precision Cosmology

DR JM NIPOK N.J.I.T.
https://orcid.org/0009-0006-3940-4450

 Feb 21 2026
Version 2.1

Abstract

Lorentz transformations are routinely used as if a single global boost could relate any cosmological source to any observer, layered on top of an FLRW background. This practice conflicts with general relativity's core structure: Lorentz invariance is strictly local, and the real universe is organized into a nested hierarchy of gravitationally bound "pockets of spacetime" — Solar System, galaxy, group, cluster, supercluster, and finally the Hubble flow. In such a universe, the physically correct transformation between a source and an observer must ascend each branch to their lowest common parent frame, evaluate the photon's path in that parent's proper spacetime, and then descend through the other branch, composing kinematic and gravitational shifts at each level. Current methods skip this step entirely, collapsing the frame tree into a single effective boost plus cosmological redshift. This paper identifies that omission as a structural gap in standard Lorentz usage and develops a frame-tree formalism that composes local $k \cdot u$ factors along the full null geodesic.

The resulting hierarchical redshift expression targets systematic effects at the 10^{-5}–10^{-4} level in $1 + z$—small enough to be ignored in everyday applications, but large enough to matter in cutting-edge precision cosmology, where such shifts can bias parameters like H_0 and w at the percent level. We argue that while the traditional Lorentz+FLRW treatment is "good enough" for most purposes, resolving current tensions and exploiting upcoming survey precision will require replacing flat boosts with frame trees rooted in the lowest common parent.

General relativity tells us two uncomfortable truths:

- Lorentz frames are local, not global.
- The universe is not a smooth soup, but a nested hierarchy of bound systems: Solar System → Galaxy → Group → Cluster → Supercluster → Hubble flow.

Every one of those levels defines its own "pocket of spacetime" with its own proper frame and bulk geodesic. A photon leaving a supernova in a cluster host does not magically leap into our Milky Way frame; it climbs out of its host galaxy and cluster, crosses their parent environment, merges into the Hubble flow, and then descends through our own supercluster, Local Group, and Galactic potentials before reaching the Solar System. In a GR-consistent picture, the transformation between source and observer must:

1. Locate the lowest common parent frame in the hierarchy (the smallest structure containing both).
2. Treat that parent's comoving frame as the proper relative spacetime for the pair.
3. Compose local Lorentz and gravitational factors step-by-step as the photon moves up to the parent and back down the other branch.

Standard practice skips step (1) entirely. It pretends there is one big, flat pocket connecting source and observer and asks a single Lorentz boost to do a job that GR clearly says is hierarchical.

The payoff of fixing this is not in rewriting undergraduate textbooks; it's in the numbers that now matter:

- Cluster potentials and multi-scale peculiar velocities naturally generate hierarchical corrections at the 10^{-5}–10^{-4} level in redshift.
- Supernova cosmology already knows that redshift systematics of this size can shift H_0 or w by $\sim 1\%$ or more.
- Weak-field GR simulations show that "tiny" relativistic terms at the 10^{-5}–10^{-3} level are exactly where Newtonian+FLRW approximations start to fray.

The frame-tree formalism developed here does three things:

- It closes the conceptual gap, aligning Lorentz practice with GR's local-frame structure.
- It provides a concrete algorithm (identify the frame tree, find the LCP, compose $k \cdot u$ factors from root to leaves) that can sit on top of existing N-body simulations and data pipelines.
- It gives precision cosmology a new handle on subtle systematics without invoking exotic physics.

If cosmology is going to argue over percent-level differences, it cannot afford to ignore a structurally missing factor at the 10^{-5}–10^{-4} level in $1 + z$. "From Flat Boosts to Frame Trees" is an invitation to replace the last global shortcut in our kinematic toolkit with a physically honest, hierarchical alternative — and to see whether some of our most persistent tensions survive that upgrade.

Introduction

Lorentz transformations were never designed to carry the full weight of modern cosmology. They were built for a world where spacetime is flat, frames are global, and a single boost can faithfully connect two laboratories in relative motion. In that regime, the mathematical elegance matches physical reality: the same Minkowski metric underlies both frames, and nothing of dynamical spacetime intrudes. But the universe we now probe at percent-level precision is not that world. It is a layered, gravitationally sculpted cosmos in which light moves not through a single homogeneous stage, but through a sequence of distinct spacetime pockets, each with its own natural frame, clock, and potential well.

Between any astrophysical source and an observer on Earth lies a pair of nested successions of frames. On the source side, a photon may spend millions of years in the proper frame of its host star and planetary system, then cross its galaxy's potential, then its group or cluster, then the larger supercluster environment, before joining the Hubble flow. On the observer side, that same photon then traverses the large-scale structure surrounding the Local Group, falls into the Milky Way, propagates through our Galactic halo and disk, and finally enters the Solar System's barycentric frame. Each segment of this journey has two defining features: a finite time and distance spent inside that pocket's spacetime, and a well-defined relative motion and potential of that pocket with respect to its parent. Treating the entire path as if it were contained in one global inertial frame, connected by a single "effective" Lorentz boost, discards all of that structure.

General relativity already tells us how to do this correctly. The frequency of light measured by any comoving observer is given by the scalar product $k \cdot u$, where k^μ is the photon's 4-momentum and u^μ is the observer's 4-velocity. The exact redshift between emission and detection is the ratio $(k \cdot u)_{\mathrm{em}}/(k \cdot u)_{\mathrm{obs}}$, with k^μ parallel-transported along the full null geodesic. This inherently path-dependent prescription is sensitive to where the photon spends its time: every region of spacetime through which it travels, and every frame in which its energy is measured, matters. In a hierarchically clustered universe, that path naturally decomposes into a chain of local Lorentz frames, each tied to a gravitationally bound structure. The physically honest way to connect source and observer is therefore to identify their lowest common parent frame in this hierarchy, work backward down each branch, and apply Lorentz and gravitational transformations pocket by pocket.

The length of time and distance that light spends in each pocket is not a mere aesthetic detail; it modulates how much that pocket's relative velocity and potential imprint on the observed redshift. A galaxy cluster's potential, for example, produces gravitational redshift signals at the level of 10^{-5}–10^{-4}, and peculiar velocities at multiple scales routinely reach $v/c \sim 10^{-3}$. Taken individually, each effect is small; taken together, along a complex, multi-pocket path, they can produce cumulative shifts in $1 + z$ of order 10^{-5}–10^{-4}. For most applications, that is negligible. But ultra-precise cosmology now lives exactly in this regime: supernova cosmology and similar probes are sensitive enough that redshift systematics at 10^{-5}–10^{-4} can bias key parameters such as the Hubble constant H_0 and the dark energy equation-of-state parameter w at the percent level.

This paper argues that the long-standing practice of modeling source–observer transformations as a single Lorentz boost plus a background FLRW stretch is no longer sufficient at the precision frontier. That approach effectively assumes that the entire source-to-observer trajectory lies in one global pocket of spacetime and that the photon's path is characterized by one relative velocity and one homogeneous expansion factor. In reality, Lorentz invariance is local, not global, and the universe's structure is hierarchical, not featureless. To remain faithful to general relativity in the era of tension-dominated cosmology, transformations must follow the underlying frame tree: identify the lowest common parent frame of source and observer, adopt that parent's comoving metric as the base "proper relative spacetime," and then work down each nested succession, applying Lorentz and gravitational shifts according to the time and distance light spends in each pocket.

The goal of this work is not to overturn the many successes of the existing Lorentz+FLRW framework, which remains entirely adequate for most purposes, but to upgrade it where it now shows strain. By replacing flat, global boosts with a hierarchical composition of local transformations anchored in a shared parent frame, one can target precisely the 10^{-5}–10^{-4} regime in redshift where present and upcoming surveys are most vulnerable to subtle systematics. The sections that follow formalize this conceptual picture, quantify the expected size of the corrections, and propose concrete modifications to the standard mathematical treatment of Lorentz transformations so that they remain the right tool—not just for "good enough" cosmology, but for ultra-precise cosmology as well.

1. Nested Structure from General Relativity and Special Relativity

1.1 General Relativity: "Follow the Leader" and Nested Comoving Frames

In general relativity, matter does not move against a fixed background; it follows geodesics of a curved spacetime determined by the combined gravitational influence of all mass–energy. Over cosmic time, this leads to a highly structured universe: matter aggregates into gravitationally bound systems whose motions are organized by deeper and deeper potential wells. Planets fall in the spacetime geometry created by their host star, stars orbit within the potential of their galaxy, galaxies move within the gravitational fields of groups and clusters, and clusters themselves stream along the filaments and sheets of the cosmic web embedded in the global FLRW expansion.

This "follow the leader" behavior naturally generates nested comoving frames of reference. At each level, there is an approximately free-falling center-of-mass frame that defines the bulk motion of that structure within its parent potential: a Solar System barycentric frame within the Milky Way; a Galactic center-of-mass frame within the Local Group; a Local Group frame within the surrounding cluster or supercluster; and, at the largest scales, a comoving frame tied to the Hubble flow. Each of these frames is singled out by general relativity as the frame in which that structure's center of mass follows (to excellent approximation) a geodesic in the parent spacetime.

Crucially, these frames form a hierarchy, not a collection of disconnected choices. The Solar System frame is nested inside the Galactic frame, which is nested inside the Local Group frame, and so on up to the cosmological comoving frame. The overlapping gravitational wells at each level—stellar, galactic, cluster-scale—do more than simply pull bodies together: they dynamically correlate their motions, causing large sets of bodies to share similar geodesics and thereby define larger and larger comoving regions.

1.2 Special Relativity: Shared Trajectories and Shared Perceptions of Space and Time

Special relativity, formulated in flat spacetime, tells us how to describe physics in an inertial frame: a frame in which freely moving bodies travel on straight lines at constant velocity, and in which the laws of physics take their simplest form. Whenever a collection of bodies shares the same bulk trajectory and relative velocity—so that they are nearly at rest with respect to one another—they can be described in a single inertial frame to high accuracy. In that frame, they share a common notion of:

- Rest: the frame in which the bulk of the system is at rest.
- Time: a proper time along the bulk geodesic that serves as the base clock for the system.

- Space: spatial slices orthogonal to that bulk motion, providing a shared spatial geometry locally.

Applied to astrophysical structures, this means that each gravitationally bound system with a well-defined center-of-mass motion admits a proper/comoving frame in which its members share the same base perception of space and time. Stars in a galaxy, for example, have small velocity dispersions around the galaxy's bulk motion; in the galaxy's rest frame, that bulk motion defines a natural time coordinate and a local Minkowski structure. The same holds for the Solar System in the Galactic frame, the Milky Way in the Local Group frame, and so on.

Because these systems are held together by overlapping gravity wells, their internal clocks and rulers are effectively entangled by the shared potential and shared motion. Inside a given level of the hierarchy, the bodies are not just near one another; they are bound into a common dynamical environment that justifies treating them as inhabiting a single "pocket" of spacetime. Within that pocket, special relativity applies with high fidelity: local experiments reveal the flat spacetime structure appropriate to the pocket's center-of-mass inertial frame.

1.3 Pockets of Spacetime and the Frame Tree

Taken together, general relativity and special relativity imply that the universe is not one global inertial frame but a tree of nested comoving pockets of spacetime:

- General relativity supplies the geometry and hierarchy: geodesic motion in curved spacetime, gravitational binding, and a sequence of larger and larger free-fall frames as one moves up through planetary systems, galaxies, groups, clusters, and superclusters into the Hubble flow.
- Special relativity supplies the local kinematic content at each node: in each of these frames, bodies that share a bulk trajectory and relative velocity field experience a common base notion of space and time, described by a local Minkowski metric.

Each node in this tree—a Solar System barycenter, a galactic center-of-mass frame, a cluster rest frame, the cosmological comoving frame—can therefore be regarded as a pocket of spacetime. Within that pocket:

- A shared bulk geodesic defines a preferred time coordinate (the system's proper time).
- A shared bulk velocity defines the local rest frame.

- Overlapping gravitational wells enforce dynamical coherence, keeping the members' motions close to that shared geodesic.

Children pockets inherit their parent's spacetime structure and add their own small Lorentz corrections. The Solar System's proper frame, for example, is nearly inertial within the Milky Way frame, which is itself nearly inertial within the Local Group frame, and so on. The result is a hierarchical frame tree in which each level is approximately inertial in its parent, but no single level can be elevated to a globally valid inertial frame across cosmological distances.

This hierarchical picture is the natural joint consequence of GR and SR. It sets the stage for the central argument of this work: that any ultra-precise transformation between a distant source and an observer must respect this frame tree. Specifically, it must locate the lowest common parent frame shared by the source and observer, adopt that parent's comoving metric as the base "proper relative spacetime," and then work down each nested succession, applying Lorentz transformations and gravitational shifts at every pocket the light traverses. Only such a construction can faithfully represent how the finite time and distance spent in each pocket influence the observed redshift at the level demanded by high-precision cosmology.

1. 4 What General Relativity and Special Relativity Predict About Spacetime Structure

Special relativity (SR) is formulated on Minkowski spacetime, a flat, non-dynamical, four-dimensional manifold with metric

$$ds^2 = -c^2dt^2 + dx^2 + dy^2 + dz^2 (1)$$

This metric is invariant under Lorentz transformations, which relate measurements between any two inertial (non-accelerating) frames moving at constant relative velocity. For a boost along the x-axis, the standard transformation is:

$$t' = \gamma\left(t - \frac{vx}{c^2}\right), x' = \gamma(x - vt), y' = y, z' = z, \gamma = (1 - v^2/c^2)^{-1/2} (2)$$

These transformations preserve the spacetime interval (Eq. 1) and ensure that the speed of light is the same in all inertial frames.

Minkowski spacetime has four key properties:

- Homogeneity: No point is special; the metric components are the same everywhere.
- Isotropy: No direction is preferred; physics looks the same in all directions.

- Flatness: The Riemann curvature tensor vanishes; parallel transport around a loop returns a vector unchanged.
- Global inertial frames: A single inertial coordinate system can, in principle, cover the entire spacetime; any two inertial frames are related by global Lorentz transformations.

In such a setting, a shared inertial frame defines a shared perception of space and time for all bodies comoving within it: they agree on distances and time intervals up to the usual Lorentz effects when compared to other frames. The critical (usually implicit) assumption behind standard Lorentz-transform applications is that this global structure exists and that a single Lorentz boost can meaningfully relate “source” and “observer,” no matter how far apart they are.

In the real universe, however, spacetime is curved, not flat, and this global assumption fails.

1.5 General Relativity: Curved Spacetime and Local Lorentz Frames

General relativity (GR) promotes spacetime from a static backdrop to a dynamic participant. Einstein’s field equations

$$G_{\mu\nu} + \Lambda g_{\mu\nu} = \frac{8\pi G}{c^4} T_{\mu\nu} (3)$$

link the geometry of spacetime directly to its matter–energy content: the metric $g_{\mu\nu}$ is not fixed a priori but is determined by the distribution and motion of mass–energy. Every mass concentration—planet, star, galaxy, cluster—warps spacetime around it, and this curvature is the mechanism of gravity, not a small correction on top of a flat stage.

This is a fundamental departure from the Newtonian picture: instead of particles moving in a rigid Euclidean space with an external gravitational force, GR describes particles following geodesics in a curved spacetime whose curvature is generated by those same particles.

Einstein’s equivalence principle states that in a sufficiently small region, one can always choose a freely falling frame in which the effects of gravity disappear and the laws of physics reduce to those of SR. Mathematically, at any point p in a curved spacetime, there exists a local Lorentz frame (local inertial frame) such that:

$$g_{\mu\nu}(p) = \eta_{\mu\nu}, \partial_\alpha g_{\mu\nu}(p) = 0 (4)$$

Within this infinitesimal neighborhood, the metric looks Minkowskian and Lorentz transformations are valid; special relativity holds locally.

However, this construction is inherently local. As Crowell illustrates by comparing freely falling observers in Los Angeles and Mumbai: each defines a local Lorentz frame in which nearby objects follow straight worldlines, but neither frame can be extended to cover both regions without reintroducing gravitational effects. This is analogous to flat maps on a sphere: each city can be mapped without distortion, but there is no single global flat map of the entire Earth.

In GR, Lorentz invariance survives only as a local symmetry acting on tangent spaces; the global symmetry group is the diffeomorphism group of smooth coordinate transformations. Local Lorentz invariance (acting on Lorentz indices) and general covariance (acting on coordinate indices) are independent. This distinction between local and global Lorentz invariance is the structural difference between SR and GR, and it has direct implications for how reference-frame transformations must be handled whenever gravity is non-negligible.

1.6 The FLRW Universe: A Preferred Cosmological Parent Frame

When Einstein's equations are solved under the assumptions of large-scale homogeneity and isotropy, the resulting spacetime is described by the Friedmann–Lemaître–Robertson–Walker (FLRW) metric:

$$ds^2 = -c^2 dt^2 + a(t)^2 \left[\frac{dr^2}{1 - kr^2} + r^2 d\Omega^2\right] \quad (5)$$

where $a(t)$ is the scale factor and $k = 0, \pm 1$ encodes spatial curvature.

The FLRW metric picks out a natural comoving frame in which the cosmic microwave background (CMB) is isotropic and ideal "Hubble-flow" galaxies sit at fixed comoving coordinates. This is a preferred frame in the cosmological sense: observers moving with respect to it see a CMB dipole.

In Minkowski spacetime, no inertial frame is preferred; the full Lorentz group relates all inertial frames on equal footing. In an FLRW universe, the comoving frame plays the role of a top-level parent frame: it defines a cosmic time coordinate and a globally expanding spatial metric. The cosmological redshift

$$1 + z_{\text{cosmo}} = \frac{a(t_0)}{a(t_e)} \quad (6)$$

is not a Doppler effect in flat spacetime; it arises from the expansion of the metric itself as photons propagate through the FLRW geometry. No single Lorentz boost can reproduce this effect, because it accumulates continuously along the photon's path through curved, time-dependent spacetime.

Already at the background level, then, GR replaces the global Lorentz symmetry of SR with a preferred cosmological comoving frame and strictly local Lorentz invariance elsewhere.

1.7 Pockets of Spacetime: Nested Comoving Frames from GR and SR

The real universe is not a smooth FLRW fluid. On top of the cosmological background, matter clusters into a hierarchy of gravitationally bound structures, each with a natural center-of-mass free-fall frame:

- Planetary systems (Solar System barycentric frame)
- Galaxies (Galactic center-of-mass frame)
- Galaxy groups (Local Group frame)
- Clusters (e.g., Virgo Cluster frame)
- Superclusters (e.g., Laniakea Supercluster frame)
- The Hubble-flow / CMB comoving frame at the largest scales

At each level, the structure's center of mass follows an approximate geodesic in the spacetime generated by its parent level. General relativity thus produces a nested sequence of comoving frames: the Solar System is nearly in free fall in the Milky Way potential; the Milky Way is nearly in free fall in the Local Group potential; the Local Group is nearly in free fall in the surrounding supercluster potential; and all are embedded in the FLRW expansion.

Special relativity then tells us what each level "feels like" internally. Whenever a set of bodies shares a common bulk trajectory and relative velocity—so that they are nearly at rest with respect to one another—they can be described in a single inertial frame. In that frame:

- Their rest is defined by the bulk motion of the structure.
- Their time is measured by a shared proper time along the bulk geodesic.
- Their space is given by local spatial slices orthogonal to that motion.

Because gravity binds them together, the overlapping potentials entangle their clocks and rulers: they share a common dynamical environment and a common local Minkowski structure. Each gravitationally bound level can therefore be regarded as a pocket of spacetime:

- General relativity supplies the bulk geodesic and embedding in curved spacetime.
- Special relativity supplies the local inertial description of space and time appropriate to that bulk motion.

These pockets form a frame tree: child pockets are approximately inertial within their parent's free-fall frame, inheriting its notion of space and time and adding small Lorentz corrections from their own internal motions.

1.8 The Lowest Common Parent Frame and Ultra-Precise Cosmology

For any astrophysical source and any observer, their respective pockets sit on different branches of this tree. A photon emitted in a source pocket:

- Spends finite time and distance in that pocket, measured in its local proper frame.
- Climbs out of the source's potential and enters parent pockets, spending time and distance in each.
- Eventually reaches the lowest common parent (LCP) frame shared with the observer—the smallest pocket that contains both.
- Then descends through the observer's parent pockets, falling into their potentials, until it arrives in the observer's local pocket.

At each step, the photon experiences:

- Kinematic Doppler shifts from the relative motion of each pocket with respect to its parent.
- Gravitational redshift/blueshift from entering and exiting each pocket's potential well.
- The background cosmological redshift from the FLRW expansion.

General relativity encodes all of this in the scalar product $k \cdot u$: the measured frequency is $E = -k_{\mu}u^{\mu}$, and the exact redshift is

$$1 + z = \frac{(k \cdot u)_{\text{emit}}}{(k \cdot u)_{\text{obs}}}$$

with k^μ parallel transported along the null geodesic. The hierarchical picture simply recognizes that, in a structured universe, this geodesic naturally decomposes into segments associated with each pocket in the frame tree, and that a physically honest transformation must:

1. Locate the lowest common parent frame of source and observer.
2. Treat that parent's comoving metric as the base "proper relative spacetime" for the pair.
3. Work down each nested succession from the parent to the source and to the observer, applying Lorentz and gravitational transformations at each pocket according to the time and distance the photon spends there.

For most purposes, the traditional practice—approximating this entire path with a single Lorentz boost plus an FLRW stretch—is "good enough," because all relevant velocities are small and potentials shallow. But at the 10^{-5}–10^{-4} level in $1 + z$, where ultra-precise cosmology now operates, the omission of the frame tree and the lowest common parent frame becomes a structurally significant error. The remainder of this paper builds on this Section 1 foundation to show how a hierarchical, shared-parent-frame treatment of Lorentz transformations can close that gap.

2. The Hole: What Is Missing from Standard Practice

2.1 The Single-Boost Approximation in Observational Cosmology

In practical cosmology, the transformation from a distant source to an observer is almost always treated as if it were a single step between two frames, lightly decorated with corrections. The observed redshift is typically written as

$$1 + z_{\text{obs}} = (1 + z_{\text{cosmo}})(1 + z_{\text{pec}}),$$

where z_{cosmo} encodes the FLRW expansion and z_{pec} collects small peculiar-velocity corrections for the source and sometimes the observer.

Type Ia supernova cosmology illustrates this clearly. In state-of-the-art analyses (e.g., Pantheon+), redshifts are converted from heliocentric to CMB frame using exact relativistic formulae, corrected using peculiar-velocity models to account for coherent bulk flows and local structure, and accompanied by error budgets that include uncertainties from large-scale flows. These steps are careful and sophisticated, but they treat gravitational redshift contributions from the full line-of-sight hierarchy—clusters, groups, filaments, superclusters—as negligible or as an effectively constant systematic.

Yet detailed studies show that a systematic redshift as small as $\Delta z \sim 10^{-5}$–10^{-4} can bias recovered cosmological parameters at the percent level. Such shifts are naturally of the same order as gravitational redshifts from residing in over- or under-dense environments and from traversing deep potentials, but current pipelines do not attempt a frame-by-frame GR treatment of these effects across all intermediate structures.

In other words, even the most precise distance–redshift work effectively approximates the full transformation as

$$\text{source pocket} \xrightarrow{\text{1 global boost+FLRW}} \text{observer pocket},$$

with the entire intervening hierarchy of nested pockets collapsed into a single effective correction term.

2.2 Local vs. Global Views of Cosmological Redshift

The theoretical literature on cosmological redshift reinforces the tension between what GR formally requires and what is practically implemented.

Bunn and Hogg showed that cosmological redshift in an FLRW universe can be represented as the accumulation of infinitesimal Doppler shifts along a photon's path. They introduce a family of comoving observers along the null geodesic, each living in a local Minkowski patch, and demonstrate that the net redshift arises as the product of many small SR Doppler factors. This construction is explicitly local: at each step, the photon passes from one local Lorentz frame to a neighboring one.

Chodorowski revisited this idea and emphasized that if one instead parallel-transports the source's four-velocity along hypersurfaces of constant cosmic time rather than along the null geodesic, the relation between transported velocity and observed redshift is not purely Dopplerian. For small redshifts, one can cleanly split the cosmological redshift into Doppler (kinematic) and gravitational components; at higher redshifts, the split remains meaningful but becomes nontrivial and depends on the detailed FLRW dynamics.

Both analyses, despite their differences, implicitly agree on a core GR fact: frequency shifts in curved spacetime are intrinsically path-dependent and must be computed by parallel transport of four-vectors through a sequence of local Lorentz frames. The redshift of a photon is not given by a single global Lorentz boost; it is the cumulative outcome of transport through curved spacetime.

Nevertheless, observational practice compresses this entire structure into the compact form $1 + z_{\text{obs}} = (1 + z_{\text{cosmo}})(1 + z_{\text{pec}})$, with at most localized gravitational effects (strong-lensing potentials, cluster wells) treated explicitly and the rest of the hierarchy ignored.

2.3 How Simulation Codes Treat Gravity and Frames

The gap between GR's local-frame picture and standard practice is even more explicit in cosmological simulations.

2.3.1 Newtonian Gravity on an FLRW Background

Flagship cosmological N-body and hydrodynamical codes—SWIFT (FLAMINGO), GADGET-4, AREPO, RAMSES, and others—adopt a common framework:

- Gravity is computed using Newtonian dynamics (Tree, TreePM, Fast Multipole + PM), with particles interacting pairwise in a single global coordinate system.
- Cosmological expansion is handled in comoving coordinates $\mathbf{r} = a(t)\mathbf{r}'$, with the FLRW background encoded by $a(t)$ and the Friedmann equations.
- Hydrodynamics is solved using SPH or mesh-based methods, also in the comoving frame.

- Reference frames: there is effectively one global comoving frame; no explicit hierarchy of local Lorentz frames is represented.

The GADGET-4 documentation is explicit that the code operates either in "Newtonian space" or in "cosmological integrations with comoving coordinates," in both cases with Newtonian gravity and no explicit GR curvature terms or frame hierarchy. SWIFT is similarly described as "a modern, highly parallel gravity and smoothed particle hydrodynamics code" on an expanding background, not a full GR solver.

2.3.2 Weak-Field GR Extensions

There are important efforts to go beyond this approximation. Codes such as gevolution and GRAMSES implement weak-field expansions of GR around an FLRW background, tracking metric perturbations and relativistic species. The "N-body gauge" program demonstrates that, in the weak-field regime, Newtonian simulations can be mapped to a relativistic description if one adopts an appropriate gauge and treats initial conditions and observables carefully.

However, even these approaches still treat metric perturbations as smooth fields on large scales, appropriate for power spectra and correlation functions. They do not implement a fully resolved lattice of nested gravitational wells and associated local Lorentz frames—the frame tree described in Section 1. Instead, they model an effectively continuous perturbation field, not the discrete sequence of pockets that a photon traverses on its way from source to observer.

A key result from these studies is that Newtonian simulations reproduce matter clustering statistics extremely well: relativistic corrections to the power spectrum are at most $\sim 10^{-5}$ at high redshift and $\sim 10^{-3}$ at $z = 0$ on near-horizon scales. This has reinforced the perception that Newtonian-plus-FLRW simulations are "good enough" for many purposes. But these performance metrics concern bulk clustering, not the detailed frame-by-frame treatment of individual photon paths and their redshifts.

2.4 Where the Hierarchy Disappears

Across theory, observation, and simulation, a consistent pattern emerges:

1. Theory (GR): Frequency shifts are formally given by $(k \cdot u)_{\text{em}}/(k \cdot u)_{\text{obs}}$, with k^μ parallel-transported along the null geodesic through a curved spacetime and u^μ defined in strictly local Lorentz frames. Decompositions into Doppler and gravitational parts are inherently path-dependent and constructed from families of local observers.

2. Observation: Redshifts are modeled as FLRW expansion plus modest peculiar-velocity terms, with limited inclusion of gravitational redshifts (e.g., cluster wells, lenses). Systematic gravitational redshifts at the 10^{-5}–10^{-4} level are acknowledged as potentially important but are not handled via a full hierarchical GR framework.
3. Simulation: Large-scale structure is evolved using Newtonian gravity in a single comoving frame on an FLRW background. Weak-field GR corrections, when included, are implemented as smooth metric perturbations, without explicitly resolving the nested local frames or propagating individual photon geodesics through the full frame tree.

In effect, the nested hierarchy of pockets of spacetime—planetary, stellar, galactic, group, cluster, supercluster, filament, Hubble flow—collapses into one or two effective frames plus a smooth background: a "source frame," an "observer frame," and an FLRW factor. The detailed, level-by-level structure that GR and SR jointly predict (Section 1) simply disappears from the calculational machinery.

2.5 Why the Approximation Was Historically Reasonable

The single-boost approximation did not arise out of carelessness. For much of the history of relativistic astrophysics and cosmology, it was a sensible engineering choice:

- On laboratory and Solar System scales, spacetime curvature is so weak that SR in a single inertial frame accurately captures experiments, and searches for Lorentz-violation see no deviations down to fractional levels of $\sim 10^{-17}$ or better.
- On galactic and cluster scales, Newtonian gravity with modest post-Newtonian corrections describes orbits, virial motions, and much of lensing very well. A Newtonian potential atop an FLRW background appears sufficient for many practical calculations.
- On cosmological scales, ΛCDM N-body simulations reproduce the galaxy distribution and matter power spectrum remarkably well, with relativistic corrections to bulk clustering at or below the 10^{-3} level.

As long as observational cosmology operated at several-percent precision, collapsing the frame tree to a single effective boost plus FLRW expansion was unlikely to be the dominant source of error. But the situation has changed.

Modern surveys routinely aim for sub-percent precision in H_0, w, and growth-rate parameters. Supernova cosmology studies now show that redshift systematics at the level of $\sim 2 \times 10^{-5}$ can shift inferred cosmological parameters by $\sim 1\%$. At the same time, persistent tensions—Hubble tension, anomalous bulk flows, CMB anomalies—suggest that subtle, previously neglected effects may now be observationally relevant.

In that context, continuing to treat the full source-to-observer transformation as a single effective boost on top of a smooth background is no longer obviously safe. The hierarchy predicted by GR and SR may now lie within our empirical reach.

2.6 The Hole, Precisely Stated

The "glaring hole" in standard Lorentz-transformation practice can now be stated with the language of Section 1:

1. GR predicts that spacetime is curved and Lorentz invariance is strictly local. Any frequency shift must be computed via parallel transport of the photon's four-momentum along its null geodesic through a sequence of local Lorentz frames.
2. The real universe is hierarchically structured into nested pockets of spacetime—bound systems with shared bulk geodesics and proper frames—forming a frame tree from planetary systems up to the FLRW comoving frame.
3. For any source–observer pair, there exists a lowest common parent (LCP) frame in this tree: the smallest pocket whose comoving frame contains both. A GR- and SR-consistent transformation must ascend each branch to this LCP, treat the LCP's comoving metric as the base proper relative spacetime, and then descend each nested succession, applying Lorentz and gravitational transformations at every pocket the light traverses.
4. Standard observational practice replaces this full hierarchical chain with a single effective transformation: an FLRW cosmological redshift plus small peculiar-velocity and occasional localized gravitational corrections, as if one global Lorentz boost could relate "the source frame" and "the observer frame."
5. Simulation codes, which underpin much of our theoretical modeling, operate with Newtonian gravity in a single comoving coordinate system on an FLRW background and do not represent the frame tree at all.
6. No widely used framework currently composes Lorentz transformations and gravitational redshifts across the full nested hierarchy of pockets between source and observer, despite this being the natural implication of GR's local-Lorentz structure and the observed layered organization of cosmic structure.

The mismatch, then, is not a minor technical oversight but a structural omission: by failing to locate the lowest common parent frame and by skipping the ascent–parent–descent through the frame tree, standard practice asks a single flat-space Lorentz boost to do a job that, in a GR + SR universe, is inherently hierarchical. The next section constructs a framework for hierarchical, nested transformations that restores this missing structure and explicitly tracks cumulative redshift/blueshift effects through the full chain of cosmic pockets.

3. A Framework for Hierarchical Lorentz Transformations

3.1 Conceptual Overview

The previous sections established three key facts:

1. Lorentz transformations are strictly local in curved spacetime, valid only within neighborhoods where the metric is approximately Minkowskian.
2. The real universe is hierarchically structured into nested “pockets of spacetime,” each with its own natural comoving frame and gravitational potential.
3. Standard practice collapses this hierarchy into a single cosmological redshift plus one or two effective boosts between a “source frame” and an “observer frame.”

The goal of this section is to construct a GR-consistent framework that restores the missing structure. The core ideas are:

- Represent each level of the hierarchy (each pocket) by a local Lorentz frame with 4-velocity $u^{\mu}_{(i)}$ and (if needed) an approximate local metric $g^{(i)}_{\mu\nu}$.
- Describe photon propagation by parallel transport of its 4-momentum k^{μ} along its null geodesic through these pockets.
- Compute the photon's energy/frequency in each pocket as $E_{(i)} = -g^{(i)}_{\mu\nu} k^{\mu} u^{\nu}_{(i)} = (k \cdot u)_{(i)}$.
- Express the total redshift between emission and observation as a product of local factors—cosmological, gravitational, and kinematic—accumulated across the full frame tree.

This preserves GR's local Lorentz invariance and curved-spacetime structure while remaining compatible with existing large-scale approximations (FLRW backgrounds, weak-field perturbations).

3.2 The General GR Redshift Formula

In general relativity, the frequency of a photon measured by an observer with 4-velocity u^{μ} is:

$$E = h\nu = -g_{\mu\nu} k^{\mu} u^{\nu} \equiv k \cdot u (7)$$

where k^{μ} is the photon's 4-momentum (tangent to its null geodesic) and $g_{\mu\nu}$ is the spacetime metric.

Consider:

- Emission at event x_{e}^{μ}, with emitter 4-velocity u_{e}^{μ}.
- Observation at event x_{o}^{μ}, with observer 4-velocity u_{o}^{μ}.

The exact GR expression for the redshift is:

$$1 + z = \frac{\nu_{\mathrm{e}}}{\nu_{\mathrm{o}}} = \frac{(k \cdot u)_{\mathrm{e}}}{(k \cdot u)_{\mathrm{o}}} (8)$$

The photon 4-momentum k^{μ} is obtained by solving the null geodesic equation:

$$k^{\nu} \nabla_{\nu} k^{\mu} = 0, k^{\mu} k_{\mu} = 0 (9)$$

and parallel transporting k^{μ} along the photon's path from emission to observation.

Equation (8) is the foundational relation. Any split into "Doppler," "gravitational," or "cosmological" pieces is a matter of how k^{μ} and the observers' 4-velocities are expressed and transported in a chosen coordinate system or tetrad basis. The hierarchical framework does not change Eq. (8); it organizes its evaluation along the frame tree.

3.3 Discrete Hierarchical Decomposition

To make the role of the frame tree explicit, discretize the photon's path into segments associated with the nested pockets it traverses.

Let the photon pass through regions

$$\mathcal{R}_0, \mathcal{R}_1, \ldots, \mathcal{R}_N, \mathcal{R}_{N+1},$$

where:

- $\mathcal{R}_0$: source pocket.
- $\mathcal{R}_{N+1}$: observer pocket.
- $\mathcal{R}_1, \ldots, \mathcal{R}_N$: intermediate pockets along the ascent–parent–descent chain (up from source to the lowest common parent, then down to observer).

Associate to each region $\mathcal{R}_i$:

- A representative event on the photon's path, $x_{(i)}^{\mu}$.

- A local metric approximation $g_{\mu\nu}^{(i)}$ (e.g., Schwarzschild or NFW potential embedded in FLRW, or a perturbed FLRW patch).
- A local comoving observer with 4-velocity $u_{(i)}^{\mu}$.

Define the photon's energy in region i as:

$$E_{(i)} = -g_{\mu\nu}^{(i)} k_{(i)}^{\mu} u_{(i)}^{\nu} = (k \cdot u)_{(i)} (10)$$

where $k_{(i)}^{\mu}$ is the photon 4-momentum at $x_{(i)}^{\mu}$, obtained by parallel transport from $x_{(i-1)}^{\mu}$ through the actual curved spacetime.

The local redshift factor between regions $\mathcal{R}_i$ and $\mathcal{R}_{i+1}$ is then:

$$1 + z_{i \to i+1} = \frac{E_{(i)}}{E_{(i+1)}} = \frac{(k \cdot u)_{(i)}}{(k \cdot u)_{(i+1)}} (11)$$

The total redshift between emitter and observer is the product:

$$1 + z_{\text{tot}} = \prod_{i=0}^{N} (1 + z_{i \to i+1}) (12)$$

This is an exact identity if the sequence of regions is fine-grained enough. In practice, one groups segments into physically meaningful layers (host galaxy potential, cluster potential, supercluster, large-scale expansion, Local Group, Milky Way, Solar System) and assigns an approximate metric and frame to each.

This discretized picture simply makes explicit what Eq. (8) already implies: the redshift is a cumulative product of local factors associated with each pocket along the photon's path.

3.4 Cosmological, Gravitational, and Kinematic Factors at Each Level

For many cosmological applications, the metric can be approximated as a perturbed FLRW spacetime in Newtonian gauge:

$$ds^2 = -(1 + 2\Psi) c^2 dt^2 + a(t)^2 (1 - 2\Phi) d\vec{x}^{\,2} (13)$$

Here Ψ and Φ are gravitational potentials sourced by inhomogeneities, and $a(t)$ is the scale factor encoding the background expansion. For non-relativistic matter, typically $\Psi \approx \Phi$, and the Newtonian potential appears in both the temporal and spatial metric components.

In this setting, the total redshift between emitter and observer is often written schematically as:

$$1 + z_{\text{tot}} \approx (1 + z_{\text{cosmo}})(1 + z_{\text{grav}})(1 + z_{\text{kin}})(14)$$

where:

- $1 + z_{\text{cosmo}} = a(t_{\text{o}})/a(t_{\text{e}})$ is the usual cosmological redshift.
- z_{grav} encodes gravitational potential differences and integrated effects (Sachs–Wolfe, Rees–Sciama).
- z_{kin} collects kinematic Doppler shifts from peculiar velocities.

The hierarchical framework refines Eq. (14) by assigning separate gravitational and kinematic contributions to each pocket in the frame tree.

3.4.1 Hierarchical Gravitational Redshift

In a static, spherically symmetric potential (e.g., Schwarzschild-like) with Newtonian potential $\Phi(r) = -GM/r$, the gravitational redshift between radii r_1 and r_2 ($r_2 > r_1$) is, to first order:

$$1 + z_{\text{grav}} \approx 1 + \frac{\Phi(r_1) - \Phi(r_2)}{c^2}(15)$$

In a hierarchical universe, each bound structure (galaxy, group, cluster, supercluster) has its own potential well. The photon's hierarchical gravitational factor becomes:

$$1 + z_{\text{grav, hier}} \approx 1 + \sum_i \frac{\Phi_{(i,\text{in})} - \Phi_{(i,\text{out})}}{c^2}(16)$$

where $\Phi_{(i,\text{in})}$ and $\Phi_{(i,\text{out})}$ are the potentials at the photon's entry and exit points for region $\mathcal{R}_i$. For regions that are weak perturbations on FLRW, these contributions can be treated as integrated potential terms along the path (generalized Sachs–Wolfe/Rees–Sciama corrections).

3.4.2 Hierarchical Kinematic (Doppler) Factors

In a perturbed FLRW universe, the 4-velocity of a comoving observer with peculiar velocity $\vec{v}$ relative to the background can be approximated such that the observed frequency relates to the emitted one via a Doppler factor:

$$\nu_{\text{obs}} \approx \nu_{\text{emit}}\, \gamma(1 - \vec{n} \cdot \vec{v}/c)(17)$$

where $\vec{n}$ is the photon's propagation direction and $\gamma \approx 1/\sqrt{1 - v^2/c^2}$.

For a hierarchy of pockets with velocities $\vec{v}_{(i)}$ relative to their parents, the net kinematic factor is the product of Doppler factors at each interface:

$$1 + z_{\text{kin, hier}} = \prod_i \gamma_{(i)}\, (1 - \vec{n}_{(i)} \cdot \vec{v}_{(i)}/c)(18)$$

Because typical peculiar velocities are $v \ll c$, one can expand to first order:

$$1 + z_{\text{kin, hier}} \approx 1 - \sum_i \vec{n}_{(i)} \cdot \frac{\vec{v}_{(i)}}{c}\,(19)$$

The directions $\vec{n}_{(i)}$ change as the photon's path is bent by lensing and by the background geometry.

3.4.3 Combined Hierarchical Formula

Combining these pieces, a schematic hierarchical generalization of Eq. (14) is:

$$1 + z_{\text{tot}} \approx \frac{a(t_{\text{o}})}{a(t_{\text{e}})}\left[1 + \sum_i \left(\frac{\Phi(i\text{,in}) - \Phi(i\text{,out})}{c^2} - \vec{n}_{(i)} \cdot \frac{\vec{v}_{(i)}}{c}\right) + \cdots \right](20)$$

with higher-order terms in $v^2/c^2, \Phi^2/c^4$, and integrated potential/time-evolution contributions (ISW/RS) along the line of sight.

Equation (20) is not a final numerical recipe; it is a structural statement: every hierarchical level contributes separately and, in principle, measurably to the total redshift, and these contributions combine multiplicatively in the precise GR sense given by Eqs. (8), (11), and (12).

3.5 Toward Implementation: A Hierarchical Transformation Algorithm

To move from theory to practice—either in analytic work or in simulations—one can cast the framework as a concrete algorithm:

1. Specify the background
   Choose an FLRW cosmology $a(t), H(t), \Omega_i$ to define the global background metric.

2. Construct the hierarchy
   For a given source–observer pair, identify the sequence of structures and pockets along the relevant geodesic: host galaxy, group, cluster, supercluster, large-scale filaments, Local Group, Milky Way, Solar System. In simulations, this can be done by tracing rays through the matter distribution and assigning them to halos or environments at each step.

3. Find the lowest common parent
   Using the frame tree, determine the lowest common parent frame (LCP) that contains both source and observer. This defines the base proper relative spacetime for the transformation.

4. Assign local metrics and frames
   For each pocket $\mathcal{R}_i$, assign an approximate local metric (e.g., Schwarzschild/NFW in FLRW, or Newtonian-gauge perturbed FLRW) and a comoving 4-velocity $u^{\mu}_{(i)}$.

5. Propagate the photon
   Numerically integrate the null geodesic equation (Eq. 9) through the sequence of regions. At each interface, parallel transport k^{μ} into the next pocket's coordinate/tetrad basis.

6. Compute local $k \cdot u$ values
   At each representative event $x^{\mu}_{(i)}$, compute $E_{(i)} = -k^{\mu} u^{(i)}_{\mu}$. The ratios $E_{(i)}/E_{(i+1)}$ give the local factors $1 + z_{i \to i+1}$ via Eq. (11).

7. Accumulate the total redshift and compare
   Multiply all local factors to obtain $1 + z_{\text{tot}}$ as in Eq. (12). Compare this to the usual single-boost + FLRW result to quantify the size and structure dependence of hierarchical corrections.

This algorithm is modular: it can be implemented as post-processing on top of existing Newtonian N-body simulations (through halo catalogs and potential fields) or integrated into relativistic codes (gevolution, GRAMSES) that already carry metric perturbations and geodesic solvers.

3.6 Summary of the Hierarchical Framework

The hierarchical Lorentz framework rests on three central claims:

1. The exact GR redshift is given by the ratio $(k \cdot u)_{\mathrm{e}}/(k \cdot u)_{\mathrm{o}}$, with k^{μ} parallel transported along the null geodesic in the actual curved spacetime.
2. In a universe with nested structure, the geodesic between emitter and observer naturally decomposes into segments associated with distinct pockets. Each segment contributes a local factor $1 + z_{i \to i+1} = (k \cdot u)_{(i)}/(k \cdot u)_{(i+1)}$, and the total redshift is their product.
3. In perturbed FLRW cosmologies, these local factors can be organized into cosmological, gravitational, and kinematic contributions at each hierarchical level, providing a controlled GR-consistent refinement of the standard single-boost approximation.

The next part of the paper can then apply this framework to concrete examples and estimate how large the hierarchical corrections are for realistic lines of sight, and whether they are sufficient to impact current cosmological tensions.

4. Conclusion: From Flat Boosts to Frame Trees

Lorentz transformations have carried cosmology a long way by riding on a useful idealization: that a single flat-space boost, layered on top of an FLRW background, can effectively relate any source to any observer. The central message of this work is not that this picture is wrong everywhere, but that it is structurally incomplete precisely where modern cosmology now cares most—at the 10^{-5}–10^{-4} level in redshift where the percent-level parameter shifts live.

General relativity and special relativity together paint a different kinematic landscape. Spacetime in the real universe is carved into a hierarchy of comoving pockets—Solar System, galaxy, group, cluster, supercluster, Hubble flow—each with its own bulk geodesic and proper frame. Special relativity says that bodies sharing that bulk trajectory and relative velocity share a base perception of space and time in that pocket. General relativity says that photons accumulate redshift not from a single global boost, but from parallel transport of their four-momentum through this entire frame tree, with local Lorentz frames valid only patch by patch.

The glaring hole in standard Lorentz practice is the absence of the lowest common parent frame and the ascent–parent–descent logic that follows from it. Current pipelines jump directly from "source frame" to "observer frame," collapsing all intermediate pockets into one effective boost plus FLRW stretch. The framework developed here fixes that by:

- Treating the shared parent's comoving metric as the proper relative spacetime for the pair.
- Discretizing the photon path into pocket-by-pocket segments.
- Composing local cosmological, gravitational, and kinematic factors along both nested branches from parent to source and from parent to observer.

For astrophysicists, this offers a two-fold opportunity:

1. Improved precision in observations.

   As surveys push toward sub-percent constraints on parameters like H_0 and w, systematic redshift errors at the 10^{-5}–10^{-4} levels are no longer negligible. A frame-tree, lowest-common-parent treatment provides a principled way to model how much time and distance light actually spends in each potential well and moving frame along its journey. It gives supernova cosmology, strong-lens time-delay measurements, standard siren analyses, and redshift-drift experiments a more faithful, GR-consistent kinematic backbone—one that can be layered atop existing data pipelines with modest additional modeling effort.

2. More realistic and testable simulations.

   Cosmological N-body and hydro codes have been extraordinarily successful using Newtonian gravity in a single comoving frame on an FLRW background. The proposed hierarchical framework does not require abandoning this foundation. Instead, it suggests a modular upgrade path: attach a frame tree and a hierarchical redshift module to existing simulations; propagate photons through halo catalogs and potential fields using local $k \cdot u$ factors; and quantify how often and by how much hierarchical corrections matter. Relativistic codes can go further and make the frame tree explicit in their treatment of metric perturbations and geodesics. In both cases, the result is a simulation output that speaks the same "local Lorentz + curved spacetime" language that GR does.

In short, flat boosts have been good enough for cosmology's first precision era. But as the field confronts stubborn tensions and designs experiments to measure redshift drift, percent-level shifts in H_0, and subtle gravitational redshift signals, it is time to let the theory's full structure inform the kinematics. Moving from flat boosts to frame trees—rooted in the lowest common parent and built from pocket-by-pocket Lorentz and gravitational factors—offers a conceptually clean, mathematically justified, and practically implementable way to do that. It asks only that we stop pretending the universe is one big inertial frame, and start letting our transformations follow the hierarchy that the universe's own gravity has already built.

Acknowledgements

The author thanks the broader cosmology and relativity communities for decades of work clarifying the interplay between local Lorentz invariance, curved spacetime, and large-scale structure, which made this synthesis possible. Discussions and published analyses on peculiar velocities, gravitational redshift in galaxy clusters, and weak-field general-relativistic N-body methods have been particularly influential in highlighting both the strengths and the limitations of current practice. The author is also grateful for ongoing efforts in high-precision supernova cosmology and large-volume simulations, whose pursuit of sub-percent accuracy strongly motivated the hierarchical framework proposed here. The author also acknowledges the assistance of the Perplexity AI assistant, whose iterative analysis and synthesis of the literature helped refine the conceptual structure and presentation of this work.

References

1. Carroll, S. M. (2004), *Spacetime and Geometry: An Introduction to General Relativity*, Addison-Wesley.
2. Misner, C. W., Thorne, K. S., & Wheeler, J. A. (1973), *Gravitation*, W. H. Freeman.
3. Bunn, E. F., & Hogg, D. W. (2009), "The kinematic origin of the cosmological redshift," *American Journal of Physics* 77, 688.
4. Chodorowski, M. J. (2011), "Is cosmological redshift a relativistic Doppler shift?," *Monthly Notices of the Royal Astronomical Society* 413, 585.
5. Adamek, J., Durrer, R., & Kunz, M. (2014), "Relativistic N-body simulations of structure formation," *Classical and Quantum Gravity* 31, 234006.
6. Josh Calcino and Tamara Davis JCAP01(2017)038 "The need for accurate redshifts in supernova cosmology," *Monthly Notices of the Royal Astronomical Society* 502, 2539.
7. Schaye, J., et al. (2023), "The FLAMINGO project: cosmological hydrodynamical simulations for large-scale structure and galaxy formation," *Monthly Notices of the Royal Astronomical Society* 520, 621.