



Graded Modal Dependent Type Theory

Benjamin Moon¹, Harley Eades III², and Dominic Orchard¹

¹ University of Kent, Canterbury, UK
{bgm4,d.a.orchard}@kent.ac.uk

² Augusta University, Augusta, USA
harley.eades@gmail.com

Abstract. Graded type theories are an emerging paradigm for augmenting the reasoning power of types with parameterizable, fine-grained analyses of program properties. There have been many such theories in recent years which equip a type theory with quantitative dataflow tracking, usually via a semiring-like structure which provides analysis on variables (often called ‘quantitative’ or ‘coeffect’ theories). We present Graded Modal Dependent Type Theory (GRTT for short), which equips a dependent type theory with a general, parameterizable analysis of the flow of data, both in and between computational terms and types. In this theory, it is possible to study, restrict, and reason about data use in programs and types, enabling, for example, parametric quantifiers and linearity to be captured in a dependent setting. We propose GRTT, study its metatheory, and explore various case studies of its use in reasoning about programs and studying other type theories. We have implemented the theory and highlight the interesting details, including showing an application of grading to optimising the type checking procedure itself.

1 Introduction

The difference between simply-typed, polymorphically-typed, and dependently-typed languages can be characterised by the *dataflow* permitted by each type theory. In each, dataflow can be enacted by *substituting* a term for occurrences of a variable in another term, the scope of which is delineated by a binder. In the simply-typed λ -calculus, data can only flow in ‘computational’ terms; computations and types are separate syntactic categories, with variables, bindings (λ), and substitution—and thus dataflow—only at the computational level. In contrast, polymorphic calculi like System F [26,52] permit dataflow within types, via type quantification (\forall), and a limited form of dataflow from computations to types, via type abstraction (Λ) and type application. Dependently-typed calculi (e.g., [14,40,41,42]) break down the barrier between computations and types further: variables are bound simultaneously in types and computations, such that data can flow both to computations and types via dependent functions (Π) and application. This pervasive dataflow enables the Curry-Howard correspondence to be leveraged for program reasoning and theorem proving [59]. However, unrestricted dataflow between computations and types can impede reasoning and can interact poorly with other type theoretic ideas.

Firstly, System F allows *parametric reasoning* and notions of representation independence [53,57], but this is lost in general in dependently-typed languages when quantifying over higher-kinded types [45] (rather than just ‘small’ types [7,36]). Furthermore, unrestricted dataflow impedes efficient compilation as compilers do not know, from the types alone, where a term is actually needed. Additional static analyses are needed to recover dataflow information for optimisation and reasoning. For example, a term shown to be used only for type checking (not flowing to the computational ‘run time’ level) can be erased [9]. Thus, dependent theories do not expose the distinction between proof relevant and irrelevant terms, requiring extensions to capture irrelevance [4,50,51]. Whilst unrestricted dataflow between computations and terms has its benefits, the permissive nature of dependent types can hide useful information. This permissiveness also interacts poorly with other type theories which seek to deliberately restrict dataflow, notably *linear types*.

Linear types allow data to be treated as a ‘resource’ which must be consumed exactly once: linearly-typed values are restricted to linear dataflow [27,58,60]. Reasoning about resourceful data has been exploited by several languages, e.g., ATS [54], Alms [56], Clean [18], Granule [46], and Linear Haskell [8]. However, linear dataflow is rare in a dependently-typed setting. Consider typing the body of the polymorphic identity function in Martin-Löf type theory:

$$a : \text{Type}, x : a \vdash x : a$$

This judgment uses a twice (typing x in the context and the subject of the judgment) and x once in the term but not at all in the type. There have been various attempts to meaningfully reconcile linear and dependent types [12,15,37,39] usually by keeping them separate, allowing types to depend only on non-linear variables. All such theories cannot distinguish variables used for computation from those used purely for type formation, which could be erased at runtime.

Recent work by McBride [43], refined by Atkey [6], generalises ideas from ‘coeffect analyses’ (variable usage analyses, like that of Petricek et al. [49]) to a dependently-typed setting to reconcile the ubiquitous flow of data in dependent types with the restricted dataflow of linearity. This approach, called Quantitative Type Theory (QTT), types the above example as:

$$a \overset{0}{:} \text{Type}, x \overset{1}{:} a \vdash x \overset{1}{:} a$$

The annotation 0 on a explains that we can use a to form a type, but we cannot, or do not, use it at the term level, thus it can be erased at runtime. The cornerstone of QTT’s approach is that dataflow of a term to the type level counts as 0 use, so arbitrary type-level use is allowed whilst still permitting quantitative analysis of computation-level dataflow. Whilst this gives a useful way to relate linear and dependent types, it cannot however reason about dataflow at the type-level (all type-level usage counts as 0). Thus, for example, QTT cannot express that a variable is used just computationally but not at all in types.

In an extended abstract, Abel proposes a generalisation of QTT to track variable use in both types and computations [2], suggesting that tracking in types

enables type checking optimisations and increased expressivity. We develop a core dependent type theory along the same lines, using the paradigm of *grading*: graded systems augment types with additional information, capturing the structure of programs [23,46]. We therefore name our approach *Graded Modal Dependent Type Theory* (GRTT for short). Our type theory is parameterised by a semiring which, like other coefficient and quantitative approaches [3,6,10,25,43,49,61], describes dataflow through a program, but in *both types and computations equally*, remedying QTT’s inability to track type-level use. We extend Abel’s initial idea by presenting a rich language, including dependent tensors, a complete metatheory, and a *graded modality* which aids the practical use of this approach (e.g., enabling functions to use components of data non-uniformly). The result is a calculus which extends the power of existing non-dependent graded languages, like Granule [46], to a dependent setting.

We begin with the definition of GRTT in Section 2, before demonstrating the power of GRTT through case studies in Section 3, where we show how to use grading to restrict GRTT terms to simply-typed reasoning, parametric reasoning (regaining universal quantification smoothly within a dependent theory), existential types, and linear types. The calculus can be instantiated to different kinds of dataflow reasoning: we show an example application to information-flow security. We then show the metatheory of GRTT in Section 4: admissibility of graded structural rules, substitution, type preservation, and strong normalisation.

We implemented a prototype language based on GRTT called **Gerty**.³ We briefly mention its syntax in Section 2.5 for use in examples. Later, Section 5 describes how the formal definition of GRTT is implemented as a bidirectional type checking algorithm, interfacing with an SMT solver to solve constraints over grades. Furthermore, Abel conjectured that a quantitative dependent theory could enable usage-based optimisation of type-checking itself [2], which would assist dependently-typed programming at scale. We validate this claim in Section 5 showing a grade-directed optimisation to **Gerty**’s type checker.

Section 6 discusses next steps for increasing the expressive power of GRTT. Full proofs and details are provided in the extended version of this paper [44].

Gerty has some similarity to Granule [46]: both are functional languages with graded types. However, Granule has a linearly typed core and no dependent types (only indexed types), thus has no need for resource tracking at the type level (type indices are not subject to tracking and their syntax is restricted).

2 GrTT: Graded Modal Dependent Type Theory

GRTT augments a standard presentation of dependent type theory with ‘grades’ (elements of a semiring) which account for how variables are used, i.e., their *dataflow*. Whilst existing work uses grades to describe usage only in computational terms (e.g. [10]), GRTT incorporates additional grades to account for how variables are used in types. We introduce here the syntax and typing, and briefly show the syntax of the implementation. Section 4 describes its metatheory.

³ <https://github.com/granule-project/gerty/releases/tag/esop2021>

2.1 Syntax

The syntax of GRTT is that of a standard Martin-Löf type theory, with the addition of a *graded modality* and grade annotations on function and tensor binders. Throughout, s and r range over grades, which are elements of a semiring $(\mathcal{R}, *, 1, +, 0)$. It is instructive to instantiate this semiring to the natural number semiring $(\mathbb{N}, \times, 1, +, 0)$, which captures the exact number of times variables are used. We appeal to this example in descriptions here.

GRTT has a single syntactic sort for computations and types:

| | | | | | | | | | | |
|-----------------|--|--|---------------------------------|-----------------------|-----------------|--|---------------|--------------|-----------|--|
| <i>(terms)</i> | $t, A, B, C ::= x$ | | $(x :_{(s,r)} A) \rightarrow B$ | | Type_l | | $\lambda x.t$ | | $t_1 t_2$ | |
| | | | | $(x :_r A) \otimes B$ | | | | (t_1, t_2) | | $\text{let } (x, y) = t_1 \text{ in } t_2$ |
| | | | | $\Box_s A$ | | | | $\Box t$ | | $\text{let } \Box x = t_1 \text{ in } t_2$ |
| <i>(levels)</i> | $l ::= 0 \mid \text{suc } l \mid l_1 \sqcup l_2$ | | | | | | | | | |

Terms include variables and a constructor for an inductive hierarchy of universes, annotated by a level l . Dependent function types are annotated with a pair of grades s and r , with s capturing how x is used in the body of the inhabiting function and r capturing how x is used in the codomain B . Dependent tensors have a single grade r , which describes how the first element is used in the typing of the second. The graded modal type operator $\Box_s A$ ‘packages’ a term and its dependencies so that values of type A can be used with grade s in the future. Graded modal types are introduced via *promotion* $\Box t$ and eliminated via $\text{let } \Box x = t_1 \text{ in } t_2$. The following sections explain the semantics of each piece of syntax with respect to its typing. We typically use A and B to connote terms used as types.

2.2 Typing Judgments, Contexts, and Grading

Typing judgments are written in either of the following two equivalent forms:

$$(\Delta \mid \sigma_1 \mid \sigma_2) \odot \Gamma \vdash t : A \qquad \left(\begin{array}{c} \Delta \\ \sigma_1 \\ \sigma_2 \end{array} \right) \odot \Gamma \vdash t : A$$

The ‘horizontal’ syntax (left) is used most often, with the equivalent ‘vertical’ form (right) used for clarity in some places. Ignoring the part to the left of \odot , typing judgments and their rules are essentially those of Martin-Löf type theory (with the addition of the modality) where Γ ranges over usual dependently-typed typing *contexts*. The left of \odot provides the grading information, where σ and Δ range over *grade vectors* and *context grade vectors* respectively, of the form:

$$\begin{array}{lll} \text{(contexts)} & \text{(grade vectors)} & \text{(context grade vectors)} \\ \Gamma ::= \emptyset \mid \Gamma, x : A & \sigma ::= \emptyset \mid \sigma, s & \Delta ::= \emptyset \mid \Delta, \sigma \end{array}$$

A grade vector σ is a vector of semiring elements, and a context vector Δ is a vector of grade vectors. We write (s_1, \dots, s_n) to denote an n -vector and likewise for context grade vectors. We omit parentheses when this would not cause ambiguity. Throughout, a comma is used to concatenate vectors and disjoint contexts, and to extend vectors with a single grade, grade vector, or typing assumption.

For a judgment $(\Delta \mid \sigma_s \mid \sigma_r) \odot \Gamma \vdash t : A$ the vectors Γ , Δ , σ_s , and σ_r are all of equal size. Given a typing assumption $y : B$ at index i in Γ , the grade $\sigma_s[i] \in \mathcal{R}$ denotes the use of y in t (the *subject* of the judgment), the grade $\sigma_r[i] \in \mathcal{R}$ denotes the use of y in A (the *subject's type*), and $\Delta[i] \in \mathcal{R}^i$ (of size i) describes how assumptions prior to y are used to form y 's type, B .

Consider the following example, which types the body of a function that takes two arguments of type a , and returns only the first:

$$\left(\begin{matrix} (0, (1), (1, 0)) \\ 0, 1, 0 \\ 1, 0, 0 \end{matrix} \right) \odot a : \text{Type}_l, x : a, y : a \vdash x : a$$

Let the context grade vector be called Δ . Then, $\Delta[0] = ()$ (empty vector) explains that there are no assumptions that are used to type a in the context, as Type_l is a closed term and the first assumption. $\Delta[1] = (1)$ explains that the first assumption a is used (grade 1) in the typing of x in the context, and $\Delta[2] = (1, 0)$, explains that a is used once in the typing of y in the context, and x is unused in the typing of y . The subject grade vector $\sigma_s = (0, 1, 0)$ explains that a is unused in the subject, x is used once, and y is unused. Finally, the subject type vector $\sigma_r = (1, 0, 0)$ explains that a appears once in the subject's type (which is just a), and x and y are unused in the formation of the subject's type.

To aid reading, recall that standard typing rules typically have the form $\text{context} \vdash \text{subject} : \text{subject-type}$, the order of which is reflected by $(\Delta \mid \sigma_s \mid \sigma_r) \odot \dots$ giving the context, subject, and subject-type grading respectively.

Well-formed Contexts The relation $\Delta \odot \Gamma \vdash$ identifies a context Γ as well-formed with respect to context grade vector Δ , defined by the following rules:

$$\frac{}{\emptyset \odot \emptyset \vdash} \text{WF}\emptyset \qquad \frac{(\Delta \mid \sigma \mid \mathbf{0}) \odot \Gamma \vdash A : \text{Type}_l}{\Delta, \sigma \odot \Gamma, x : A \vdash} \text{WFEXT}$$

Unlike typing, well-formedness does not need to include subject and subject-type grade vectors, as it considers only the well-formedness of the assumptions in a context with respect to prior assumptions in the context. The $\text{WF}\emptyset$ rule states that the empty context is well-formed with an empty context grade vector as there are no assumptions to account for. The WFEXT rule states that given A is a type under the assumptions in Γ , with σ accounting for the usage of Γ variables in A , and Δ accounting for usage within Γ , then we can form the well-formed context $\Gamma, x : A$ by extending Δ with σ to account for the usage of A in forming the context. The notation $\mathbf{0}$ denotes a vector for which each element is the semiring 0. Note that the well-formedness $\Delta \odot \Gamma \vdash$ is inherent from the premise of WFEXT due to the following lemma:

Lemma 1 (Typing contexts are well-formed). *If $(\Delta \mid \sigma_1 \mid \sigma_2) \odot \Gamma \vdash t : A$ then $\Delta \odot \Gamma \vdash$.*

2.3 Typing Rules

We examine the typing rules of GRTT one at a time.

Variables are introduced as follows:

$$\frac{(\Delta_1, \sigma, \Delta_2) \odot \Gamma_1, x : A, \Gamma_2 \vdash \quad |\Delta_1| = |\Gamma_1|}{(\Delta_1, \sigma, \Delta_2 \mid \mathbf{0}^{|\Delta_1|}, \mathbf{1}, \mathbf{0} \mid \sigma, \mathbf{0}, \mathbf{0}) \odot \Gamma_1, x : A, \Gamma_2 \vdash x : A} \text{VAR}$$

The premise identifies $\Gamma_1, x : A, \Gamma_2$ as well-formed under the context grade vector $\Delta_1, \sigma, \Delta_2$. By the size condition $|\Delta_1| = |\Gamma_1|$, we are able to identify σ as capturing the usage of the variables Γ_1 in forming A . This information is used in the conclusion, capturing type-level variable usage as $\sigma, \mathbf{0}, \mathbf{0}$, which describes that Γ_1 is used according to σ in the subject's type (A), and that the x and the variables of Γ_2 are used with grade 0. For subject usage, we annotate the first zero vector with a size $|\Delta_1|$, allowing us to single out x as being the only assumption used with grade 1 in the subject; all other assumptions are used with grade 0.

For example, typing the body of the polymorphic identity ends with VAR:

$$\frac{\dots \quad \frac{(((), (1)) \odot a : \mathbf{Type}, x : a \vdash \quad \text{WFEXT} \quad |((())| = |a : \mathbf{Type}|}{(((), (1)) \mid \mathbf{0}, \mathbf{1} \mid \mathbf{1}, \mathbf{0}) \odot a : \mathbf{Type}, x : a \vdash x : a} \text{VAR}}{\dots} \text{VAR}$$

The premise implies that $(((), (1, 0)) \odot a : \mathbf{Type} \vdash a : \mathbf{Type}$ by the following lemma:

Lemma 2 (Typing an assumption in a well-formed context). *If $\Delta_1, \sigma, \Delta_2 \odot \Gamma_1, x : A, \Gamma_2 \vdash$ with $|\Delta_1| = |\Gamma_1|$, then $(\Delta_1 \mid \sigma \mid \mathbf{0}) \odot \Gamma_1 \vdash A : \mathbf{Type}_l$ for some l .*

In the conclusion of VAR, the typing $(((), (1, 0)) \odot a : \mathbf{Type} \vdash a : \mathbf{Type}$ is ‘distributed’ to the typing of x in the context and to the formation of the subject's type. Thus subject grade $(0, 1)$ corresponds to the absence of a from the subject and the presence of x , and subject-type grade $(1, 0)$ corresponds to the presence of a in the subject's type (a), and the absence of x .

Typing universes are formed as follows:

$$\frac{\Delta \odot \Gamma \vdash}{(\Delta \mid \mathbf{0} \mid \mathbf{0}) \odot \Gamma \vdash \mathbf{Type}_l : \mathbf{Type}_{\text{succ } l}} \text{Type}$$

We use an inductive hierarchy of universes [47] with ordering $<$ such that $l < \text{succ } l$. Universes can be formed under any well-formed context, with every assumption graded with 0 subject and subject-type use, capturing the absence of any assumptions from the universes, which are closed forms.

Functions Function types $(x :_{(s,r)} A) \rightarrow B$ are annotated with two grades: explaining that x is used with grade s in the body of the inhabiting function and with grade r in B . Function types have the following formation rule:

$$\frac{(\Delta \mid \sigma_1 \mid \mathbf{0}) \odot \Gamma \vdash A : \mathbf{Type}_{l_1} \quad (\Delta, \sigma_1 \mid \sigma_2, r \mid \mathbf{0}) \odot \Gamma, x : A \vdash B : \mathbf{Type}_{l_2}}{(\Delta \mid \sigma_1 + \sigma_2 \mid \mathbf{0}) \odot \Gamma \vdash (x :_{(s,r)} A) \rightarrow B : \mathbf{Type}_{l_1 \sqcup l_2}} \rightarrow$$

The usage of the dependencies of A and B (excepting x) are given by σ_1 and σ_2 in the premises (in the ‘subject’ position) which are combined as $\sigma_1 + \sigma_2$ (via

pointwise vector addition using the $+$ of the semiring), which serves to *contract* the dependencies of the two types. The usage of x in B is captured by r , and then internalised to the binder in the conclusion of the rule. An arbitrary grade for s is allowed here as there is no information on how x is used in an inhabiting function body. Function terms are then typed by the following rule:

$$\frac{(\Delta, \sigma_1 \mid \sigma_3, r \mid \mathbf{0}) \odot \Gamma, x : A \vdash B : \mathbf{Type}_l \quad (\Delta, \sigma_1 \mid \sigma_2, s \mid \sigma_3, r) \odot \Gamma, x : A \vdash t : B}{(\Delta \mid \sigma_2 \mid \sigma_1 + \sigma_3) \odot \Gamma \vdash \lambda x.t : (x :_{(s,r)} A) \rightarrow B} \lambda_i$$

The second premise types the body of the λ -term, showing that s captures the usage of x in t and r captures the usage of x in B ; the subject and subject-type grades of x are then internalised as annotations on the function type's binder.

Dependent functions are eliminated through application:

$$\frac{(\Delta, \sigma_1 \mid \sigma_3, r \mid \mathbf{0}) \odot \Gamma, x : A \vdash B : \mathbf{Type}_l \quad (\Delta \mid \sigma_2 \mid \sigma_1 + \sigma_3) \odot \Gamma \vdash t_1 : (x :_{(s,r)} A) \rightarrow B \quad (\Delta \mid \sigma_4 \mid \sigma_1) \odot \Gamma \vdash t_2 : A}{(\Delta \mid \sigma_2 + s * \sigma_4 \mid \sigma_3 + r * \sigma_4) \odot \Gamma \vdash t_1 t_2 : [t_2/x]B} \lambda_e$$

where $*$ is the scalar multiplication of a vector, using the semiring multiplication. Given a function t_1 which uses its parameter with grade s to compute and with grade r in the typing of the result, we can apply it to a term t_2 , provided that we have the resources required to form t_2 scaled by s at the subject level and by r at the subject-type level, since t_2 is substituted into the return type B . This scaling behaviour is akin to that used in coeffect calculi [25,49], QTT [6,43] and Linear Haskell [8], but scalar multiplication happens here at both the subject and subject-type level. The use of variables in A is accounted for by σ_1 as explained in the third premise, but these usages are not present in the resulting application since A no longer appears in the types or the terms.

Consider the constant function $\lambda x.\lambda y.x : (x :_{(1,0)} A) \rightarrow (y :_{(0,0)} B) \rightarrow A$ (for some A and B). Here the resources required for the second parameter will always be scaled by 0, which is absorbing, meaning that anything passed as the second argument has 0 subject and subject-type use. This example begins to show some of the power of grading—the grades capture the program structure at all levels.

Tensors The rule for forming dependent tensor types is as follows:

$$\frac{(\Delta \mid \sigma_1 \mid \mathbf{0}) \odot \Gamma \vdash A : \mathbf{Type}_l \quad (\Delta, \sigma_1 \mid \sigma_2, r \mid \mathbf{0}) \odot \Gamma, x : A \vdash B : \mathbf{Type}_l}{(\Delta \mid \sigma_1 + \sigma_2 \mid \mathbf{0}) \odot \Gamma \vdash (x :_r A) \otimes B : \mathbf{Type}_l} \otimes$$

This rule is almost identical to function type formation \rightarrow but with only a single grade r on the binder, since x is only bound in B (the type of the second component), and not computationally. For ‘quantitative’ semirings, where 0 really means unused (see Section 3), $(x :_0 A) \otimes B$ is then a product $A \times B$.

Dependent tensors are introduced as follows:

$$\frac{(\Delta, \sigma_1 \mid \sigma_3, r \mid \mathbf{0}) \odot \Gamma, x : A \vdash B : \mathbf{Type}_l \quad (\Delta \mid \sigma_2 \mid \sigma_1) \odot \Gamma \vdash t_1 : A \quad (\Delta \mid \sigma_4 \mid \sigma_3 + r * \sigma_2) \odot \Gamma \vdash t_2 : [t_1/x]B}{(\Delta \mid \sigma_2 + \sigma_4 \mid \sigma_1 + \sigma_3) \odot \Gamma \vdash (t_1, t_2) : (x :_r A) \otimes B} \otimes_i$$

In the typing premise for t_2 , occurrences of x are replaced with t_1 in the type, ensuring that the type of the second component (t_2) is calculated using the first component (t_1). The resources for t_1 in this substitution are scaled by r , accounting for the existing usage of x in B . In the conclusion, we see the resources for the two components (and their types) combined via the semiring addition.

Finally, tensors are eliminated with the following rule:

$$\frac{(\Delta \mid \sigma_3 \mid \sigma_1 + \sigma_2) \odot \Gamma \vdash t_1 : (x :_r A) \otimes B \quad (\Delta, (\sigma_1 + \sigma_2) \mid \sigma_5, r' \mid \mathbf{0}) \odot \Gamma, z : (x :_r A) \otimes B \vdash C : \mathbf{Type}_l}{(\Delta, \sigma_1, (\sigma_2, r) \mid \sigma_4, s, s \mid \sigma_5, r', r') \odot \Gamma, x : A, y : B \vdash t_2 : [(x, y)/z]C} \otimes_e \quad \frac{(\Delta \mid \sigma_4 + s * \sigma_3 \mid \sigma_5 + r' * \sigma_3) \odot \Gamma \vdash \text{let } (x, y) = t_1 \text{ in } t_2 : [t_1/z]C}{\otimes_e}$$

As this is a dependent eliminator, we allow the result type C to depend upon the value of the tensor as a whole, bound as z in the second premise with grade r' , into which is substituted our actual tensor term t_1 in the conclusion.

Eliminating a tensor (t_1) requires that we consider each component (x and y) is used with the same grade s in the resulting expression t_2 , and that we scale the resources of t_1 by s . This is because we cannot inspect t_1 itself, and semiring addition is not injective (preventing us from splitting the grades required to form t_1). This prevents forming certain functions (e.g., projections) under some semirings, but this can be overcome by the introduction of *graded modalities*.

Graded Modality Graded binders alone do not allow different parts of a value to be used differently, e.g., computing the length of a list ignores the elements, projecting from a pair discards one component. We therefore introduce a *graded modality* (à la [10,46]) allowing us to capture the notion of local inspection on data and internalising usage information into types. A type $\square_s A$ denotes terms of type A that are used with grade s . Type formation and introduction rules are:

$$\frac{(\Delta \mid \sigma \mid \mathbf{0}) \odot \Gamma \vdash A : \mathbf{Type}_l}{(\Delta \mid \sigma \mid \mathbf{0}) \odot \Gamma \vdash \square_s A : \mathbf{Type}_l} \square \quad \frac{(\Delta \mid \sigma_1 \mid \sigma_2) \odot \Gamma \vdash t : A}{(\Delta \mid s * \sigma_1 \mid \sigma_2) \odot \Gamma \vdash \square t : \square_s A} \square_i$$

To form a term of type $\square_s A$, we ‘promote’ a term t of type A by requiring that we can use the resources used to form t (σ_1) according to grade s . This ‘promotion’ resembles that of other graded modal systems (e.g., [3,10,23,46]), but the elimination needs to also account for type usage due to dependent elimination.

We can see promotion \square_i as capturing t for later use according to grade s . Thus, when eliminating a term of type $\square_s A$, we must consider how the ‘unboxed’ term is used with grade s , as per the following dependent eliminator:

$$\frac{(\Delta, \sigma_2 \mid \sigma_4, r \mid \mathbf{0}) \odot \Gamma, z : \square_s A \vdash B : \mathbf{Type}_l \quad (\Delta \mid \sigma_1 \mid \sigma_2) \odot \Gamma \vdash t_1 : \square_s A \quad (\Delta, \sigma_2 \mid \sigma_3, s \mid \sigma_4, (s * r)) \odot \Gamma, x : A \vdash t_2 : [\square x/z]B}{(\Delta \mid \sigma_1 + \sigma_3 \mid \sigma_4 + r * \sigma_1) \odot \Gamma \vdash \text{let } \square x = t_1 \text{ in } t_2 : [t_1/z]B} \square_e$$

This rule can be understood as a kind of ‘cut’, connecting a ‘capability’ to use a term of type A according to grade s with the requirement that $x : A$ is used according to grade s as a dependency of t_2 . Since we are in a dependently-typed

setting, we also substitute t_1 into the type level such that B can depend on t_1 according to grade r which then causes the dependencies of t_1 (σ_1) to be scaled-up by r and added to the subject-type grading.

Equality, Conversion, and Subtyping A key part of dependent type theories is a notion of term equality and type conversion [33]. GRTT term equality is via judgments $(\Delta \mid \sigma_1 \mid \sigma_2) \odot \Gamma \vdash t_1 = t_2 : A$ equating terms t_1 and t_2 of type A . Equality includes full congruences as well as $\beta\eta$ -equality for functions, tensors, and graded modalities, of which the latter are:

$$\frac{(\Delta, \sigma_2 \mid \sigma_4, r \mid \mathbf{0}) \odot \Gamma, z : \square_s A \vdash B : \mathbf{Type}_i \quad (\Delta \mid \sigma_1 \mid \sigma_2) \odot \Gamma \vdash t_1 : A \quad (\Delta, \sigma_2 \mid \sigma_3, s \mid \sigma_4, (s * r)) \odot \Gamma, x : A \vdash t_2 : [\square x/z]B}{(\Delta \mid \sigma_3 + s * \sigma_1 \mid \sigma_4 + s * r * \sigma_1) \odot \Gamma \vdash (\text{let } \square x = \square t_1 \text{ in } t_2) = [t_1/x]t_2 : [\square t_1/z]B} \text{EQ}_{\square_c}$$

$$\frac{(\Delta \mid \sigma_1 \mid \sigma_2) \odot \Gamma \vdash t : \square_s A}{(\Delta \mid \sigma_1 \mid \sigma_2) \odot \Gamma \vdash t = (\text{let } \square x = t \text{ in } \square x) : \square_s A} \text{EQ}_{\square_u}$$

A subtyping relation $((\Delta \mid \sigma) \odot \Gamma \vdash A \leq B)$ subsumes equality, adding ordering of universe levels. *Type conversion* allows re-typing terms based on the judgment:

$$\frac{(\Delta \mid \sigma_1 \mid \sigma_2) \odot \Gamma \vdash t : A \quad (\Delta \mid \sigma_2) \odot \Gamma \vdash A \leq B}{(\Delta \mid \sigma_1 \mid \sigma_2) \odot \Gamma \vdash t : B} \text{CONV}$$

The full rules for equality and subtyping are in this paper's extended version [44].

2.4 Operational Semantics

As with other graded modal calculi (e.g., [3,10,23]), the core calculus of GRTT has a Call-by-Name small-step operational semantics with reductions $t \rightsquigarrow t'$. The rules are standard, with the addition of the β -rule for the graded modality:

$$\text{let } \square x = \square t_1 \text{ in } t_2 \rightsquigarrow [t_1/x]t_2 \quad (\beta\square)$$

Type preservation and normalisation are considered in Section 4.

2.5 Implementation and Examples

To explore our theory, we provide an implementation, **Gerty**. Section 5 describes how the declarative definition of the type theory is implemented as a bidirectional type checking algorithm. We briefly mention the syntax here for use in later examples. The following is the polymorphic identity function in **Gerty**:

```
id : (a : (.0, .2) Type 0) -> (x : (.1, .0) a) -> a
id = \a -> \x -> x
```

The syntax resembles the theory, where grading terms $.n$ are syntactic sugar for a unary encoding of grades in terms of 0 and repeated addition of 1, e.g., $.2 = (.0 + .1) + .1$. This syntax can be used for grade terms of any semiring, which can be resolved to particular built-in semirings at other points of type checking.

The following shows first projection on (non-dependent) pairs, using the graded modality (at grade 0 here) to give fine-grained usage on compound data:

```
fst : (a : (.0, .2) Type 0) (b : (.0, .1) Type 0) -> <a * [.0] b> -> a
fst = \a b p -> case p of <x, y> -> let [z] = y in x
```

The implementation adds various built-in semirings, some syntactic sugar, and extras such as: a singleton *unit* type, extensions of the theory to semirings with a pre-ordering (discussed further in Section 6), and some implicit resolution. Anywhere a grade is expected, an underscore can be supplied to indicate that **Gerty** should try to resolve the grade implicitly. Grades may also be omitted from binders (see above in `fst`), in which case they are treated as implicits. Currently, implicits are handled by generating existentially quantified grade variables, and using SMT to solve the necessary constraints (see Section 5).

So far we have considered the natural numbers semiring providing an analysis of usage. We come back to this and similar examples in Section 3. To show another kind of example, we consider a lattice semiring of privacy levels (appearing elsewhere [3,23,46]) which enforces information-flow control, akin to DCC [1]. Differently to DCC, dataflow is tracked through variable dependencies, rather than through the results of computations in the monadic style of DCC.

Definition 1. [Security levels] Let $\mathcal{R} = \text{Lo} \leq \text{Hi}$ be a set of labels with $0 = \text{Hi}$ and $1 = \text{Lo}$, semiring addition as the meet and multiplication as join. Here, $1 = \text{Lo}$ treats the base notion of dataflow as being in the low security (public) domain. Variables graded with Hi must then be unused, or guarded by a graded modality. This semiring is primitive in **Gerty**; we can express the following example:

```
idLo : (a : (.0, .2) Type 0) -> (x : (Lo, Hi) a) -> a
idLo = \a -> \x -> x
-- The following is rejected as ill-typed
leak : (a : (.0, .2) Type 0) -> (x : (Hi, Hi) a) -> a
leak = \a -> \x -> idLo a x
```

The first definition is well-typed, but the second yields a typing error originating from the application in its body:

```
At subject stage got the following mismatched grades:
For 'x' expected Hi but got .1
```

where grade 1 is Lo here. Thus we can use this abstract label semiring as a way of restricting flow of data between regions (*cf.* region typing systems [31,55]). Note that the ordering is not leveraged here other than in the lattice operations.

3 Case Studies

We now demonstrate GRTT via several cases studies that focus the reasoning power of dependent types via grading. Since grading in GRTT serves to explain dataflow, we can characterise subsets of GRTT that correspond to various type theories. We demonstrate the approach with simple types, parametric polymorphism, and linearity. In each case study, we restrict GRTT to a subset by *a*

characterisation of the grades, rather than by, say, placing detailed syntactic restrictions or employing meta-level operations or predicates that restrict syntax (as one might do for example to map a subset of Martin-Löf type theory into the simply-typed λ -calculus by restriction to closed types, requiring deep inspection of type terms). Since this restriction is only on grades, we can harness the specific reasoning power of particular calculi from within the language itself, simply by specifications on grades. In the context of an implementation like **Gerty**, this amounts to using type signatures to restrict dataflow.

This section shows the power of tracking dataflow in types via grades, going beyond QTT [6] and GRAD [13]. For ‘quantitative’ semirings, a 0 type-grade means that we can recover simply-typed reasoning (Section 3.3) and distinguish computational functions from type-parameter functions for parametric reasoning (Section 3.4), embedding a grade-restricted subset of GRTT into System F.

Section 5 returns to a case study that builds on the implementation.

3.1 Recovering Martin-Löf Type Theory

When the semiring parameterising GRTT is the singleton semiring (i.e., any semiring where $1 = 0$), we have an isomorphism $\square_r A \cong A$, and grade annotations become redundant, as all grades are equal. All vectors and grades on binders may then be omitted, and we can write typing judgments as $\Gamma \vdash t : A$, giving rise to a standard Martin-Löf type theory as a special case of GRTT.

3.2 Determining Usage via Quantitative Semirings

Unlike existing systems, we can use the fine-grained grading to *guarantee* the relevance or irrelevance of assumptions in types. To do this we must consider a subset of semirings $(\mathcal{R}, *, 1, +, 0)$ called *quantitative* semirings, satisfying:

$$\begin{aligned} & \text{(zero-unique)} \quad 1 \neq 0; \\ & \text{(positivity)} \quad \forall r, s. r + s = 0 \implies r = 0 \wedge s = 0; \\ & \text{(zero-product)} \quad \forall r, s. r * s = 0 \implies r = 0 \vee s = 0. \end{aligned}$$

These axioms⁴ ensure that a 0-grade in a quantitative semiring represents irrelevant variable use. This notion has recently been proved for computational use by Choudhury et al. [13] via a heap-based semantics for grading (on computations) and the same result applies here. Conversely, in a quantitative semiring any grade other than 0 denotes relevance. From this, we can *directly* encode non-dependent tensors and arrows: in $(x :_0 A) \otimes B$ the grade 0 captures that x cannot have any computational content in B , and likewise for $(x :_{(s,0)} A) \rightarrow B$ the grade 0 explains that x cannot have any computational content in B , but may have computational use according to s in the inhabiting function. Thus,

⁴ Atkey requires *positivity* and *zero-product* for all semirings parameterising QTT [6] (as does Abel [2]). Atkey imposes this for admissibility of substitution. We need not place this restriction on GRTT to have substitution in general (Sec. 4.1).

the grade 0 here describes that elimination forms *cannot* ever inspect the variable during normalisation. Additionally, quantitative semirings can be used for encoding simply-typed and polymorphic reasoning.

Example 1. Some quantitative semirings are:

- (*Exact usage*) $(\mathbb{N}, \times, 1, +, 0)$;
- (*0-1*) The semiring over $\mathcal{R} = \{0, 1\}$ with $1 + 1 = 1$ which describes relevant *vs.* irrelevant dependencies, but no further information.
- (*None-One-Tons* [43]) The semiring on $\mathcal{R} = \{0, 1, \infty\}$ is more fine-grained than 0-1, where ∞ represents more than 1 usage, with $1 + 1 = \infty = 1 + \infty$.

3.3 Simply-typed Reasoning

As discussed in Section 1, the simply-typed λ -calculus (STLC) can be distinguished from dependently-typed calculi via the restriction of dataflow: in simple types, data can only flow at the computational level, with no dataflow within, into, or from types. We can thus view a GRTT function as simply typed when its variable is irrelevant in the type, e.g., $(x :_{(s,0)} A) \rightarrow B$ for quantitative semirings. We define a subset of GRTT restricted to simply-typed reasoning:

Definition 2. [Simply-typed GRTT] For a quantitative semiring, the following predicate $\text{STLC}(-)$ determines a subset of simply-typed GRTT programs:

$$\text{STLC}((\emptyset \mid \emptyset \mid \emptyset) \odot \emptyset \vdash t : A)$$

$$\text{STLC}((\Delta \mid \sigma_1 \mid \sigma_2) \odot \Gamma \vdash t : A) \implies \text{STLC}((\Delta, \mathbf{0} \mid \sigma_1, s \mid \sigma_2, \mathbf{0}) \odot \Gamma, x : B \vdash t : A)$$

That is, all subject-type grades are 0 (thus function types are of the form $(x :_{(s,0)} A) \rightarrow B$). A similar predicate is defined on well-formed contexts (elided), restricting context grades of well-formed contexts to only zero grading vectors.

Under the restriction of Definition 2, a subset of GRTT terms embeds into the simply-typed λ -calculus in a sound and complete way. Since STLC does not have a notion of tensor or modality, this is omitted from the encoding:

$$\llbracket x \rrbracket = x \quad \llbracket \lambda x.t \rrbracket = \lambda x.\llbracket t \rrbracket \quad \llbracket t_1 t_2 \rrbracket = \llbracket t_1 \rrbracket \llbracket t_2 \rrbracket \quad \llbracket (x :_{(s,0)} A) \rightarrow B \rrbracket_\tau = \llbracket A \rrbracket_\tau \rightarrow \llbracket B \rrbracket_\tau$$

Variable contexts of GRTT are interpreted by point-wise applying $\llbracket - \rrbracket_\tau$ to typing assumptions. We then get the following preservation of typing into the simply-typed λ -calculus, and soundness and completeness of this encoding:

Lemma 3 (Soundness of typing). *Given a derivation of $(\Delta \mid \sigma_1 \mid \sigma_2) \odot \Gamma \vdash t : A$ such that $\text{STLC}((\Delta \mid \sigma_1 \mid \sigma_2) \odot \Gamma \vdash t : A)$ then $\llbracket \Gamma \rrbracket_\tau \vdash \llbracket t \rrbracket : \llbracket A \rrbracket_\tau$ in STLC.*

Theorem 1 (Soundness and completeness of the embedding). *Given $\text{STLC}((\Delta \mid \sigma_1 \mid \sigma_2) \odot \Gamma \vdash t : A)$ and $\llbracket (\Delta \mid \sigma_1 \mid \sigma_2) \odot \Gamma \vdash t : A \rrbracket$ then for CBN reduction $\rightsquigarrow^{\text{STLC}}$ in simply-typed λ -calculus:*

$$\begin{aligned} & \text{(soundness)} \quad \forall t'. \text{ if } t \rightsquigarrow t' \quad \text{then } \llbracket t \rrbracket \rightsquigarrow^{\text{STLC}} \llbracket t' \rrbracket \\ & \text{(completeness)} \quad \forall t_a. \text{ if } \llbracket t \rrbracket \rightsquigarrow^{\text{STLC}} t_a \text{ then } \exists t'. t \rightsquigarrow t' \wedge \llbracket t' \rrbracket \equiv_{\beta\eta} t_a \end{aligned}$$

Thus, we capture simply-typed reasoning just by restricting type grades to 0 for quantitative semirings. We consider quantitative semirings again for parametric reasoning, but first recall issues with parametricity and dependent types.

3.4 Recovering Parametricity via Grading

One powerful feature of grading in a dependent type setting is the ability to recover parametricity from dependent function types. Consider the following type of functions in System F (we borrow this example from Nuyts et al. [45]):

$$\text{RI } A B \triangleq \forall \gamma. (\gamma \rightarrow A) \rightarrow (\gamma \rightarrow B)$$

Due to parametricity, we get the following notion of *representation independence* in System F: for a function $f : \text{RI } A B$, some type γ' , and terms $h : \gamma' \rightarrow A$ and $c : \gamma'$, then we know that f can only use c by applying $h c$. Subsequently, $\text{RI } A B \cong A \rightarrow B$ by parametricity [52], defined uniquely as:

$$\begin{array}{ll} \text{iso} : \text{RI } A B \rightarrow (A \rightarrow B) & \text{iso}^{-1} : (A \rightarrow B) \rightarrow \text{RI } A B \\ \text{iso } f = f A (\text{id } A) & \text{iso}^{-1} g = \Lambda \gamma. \lambda h. \lambda (c : \gamma). g(h c) \end{array}$$

In a dependently-typed language, one might seek to replace System F's universal quantifier with Π -types, i.e.

$$\text{RI}' A B \triangleq (\gamma : \text{Type}) \rightarrow (\gamma \rightarrow A) \rightarrow (\gamma \rightarrow B)$$

However, we can no longer reason parametrically about the inhabitants of such types (we cannot prove that $\text{RI}' A B \cong A \rightarrow B$) as the free interaction of types and computational terms allows us to give the following non-parametric element of $\text{RI}' A B$ over ‘large’ type instances:

$$\text{leak} = \lambda \gamma. \lambda h. \lambda c. \gamma : \text{RI}' A \text{ Type}$$

Instead of applying $h c$, the above “leaks” the type parameter γ . GRTT can recover universal quantification, and hence parametric reasoning, by using grading to restrict the data-flow capabilities of a Π -type. We can refine representation independence to the following:

$$\text{RI}'' A B \triangleq (\gamma :_{(0,2)} \text{Type}) \rightarrow (h :_{(s_1,0)} (x :_{(s_2,0)} \gamma) \rightarrow A) \rightarrow (c :_{(s_3,0)} \gamma) \rightarrow B$$

for some grades s_1 , s_2 , and s_3 , and with shorthand $2 = 1 + 1$.

If we look at the definition of *leak* above, we see that γ is used in the body of the function and thus requires usage 1, so *leak* cannot inhabit $\text{RI}'' A \text{ Type}$. Instead, *leak* would be typed differently as:

$$\text{leak} : (\gamma :_{(1,2)} \text{Type}) \rightarrow (h :_{(0,0)} (x :_{(s,0)} \gamma) \rightarrow A) \rightarrow (c :_{(0,0)} \gamma) \rightarrow \text{Type}$$

The problematic behaviour (that the type parameter γ is returned by the inner function) is exposed by the subject grade 1 on the binder of γ . We can thus define a graded universal quantification from a graded Π -typed:

$$\forall_r (\gamma : A). B \triangleq (\gamma :_{(0,r)} A) \rightarrow B \tag{1}$$

This denotes that the type parameter γ can appear freely in B described by grade r , but is irrelevant in the body of any corresponding λ -abstraction. This is akin to the work of Nuyts et al. who develop a system with several modalities for regaining parametricity within a dependent type theory [45]. Note however that parametricity is recovered for us here as one of many possible options coming from systematically specialising the grading.

Capturing Existential Types With the ability to capture universal quantifier, we can similarly define existentials (allowing, e.g., abstraction [11]). We define the existential type via a Church-encoding as follows:

$$\exists_r(x : A).B \triangleq \forall_2(C : \mathbf{Type}_l).(f :_{(1,0)} \forall_r(x : A).(b :_{(s,0)} B) \rightarrow C) \rightarrow C$$

Embedding into Stratified System F We show that parametricity is regained here (and thus eqn. (1) really behaves as a universal quantifier and not a general Π -type) by showing that we can embed a subset of GRTT into System F, based solely on a classification of the grades. We follow a similar approach to Section 3.3 for simply-typed reasoning but rather than defining a purely syntactic encoding (and then proving it type sound) our encoding is type directed since we embed GRTT functions of type $(x :_{(0,r)} \mathbf{Type}_l) \rightarrow B$ as universal types in System F with corresponding type abstractions (λ) as their inhabitants. Since GRTT employs a predicative hierarchy of universes, we target Stratified System F (hereafter SSF) since it includes the analogous inductive hierarchy of kinds [38]. We use the formulation of Eades and Stump [21] with terms t_s and types T :

$$t_s ::= x \mid \lambda(x : T).t_s \mid t_s t'_s \mid \Lambda(X : K).t_s \mid t_s [T] \quad T ::= X \mid T \rightarrow T' \mid \forall(X : K).T$$

with kinds $K ::= \star_l$ where $l \in \mathbb{N}$ providing the stratified kind hierarchy. Capitalised variables X are System F type variables and $t_s [T]$ is type application. Contexts may contain both type and computational variables, and so free-variable type assumptions may have dependencies, akin to dependent type systems. Kinding is via judgments $\Gamma \vdash T : \star_l$ and typing via $\Gamma \vdash t : T$.

We define a type directed encoding on a subset of GRTT typing derivations characterised by the following predicate:

$$\begin{aligned} & \text{SSF}(\emptyset \mid \emptyset \mid \emptyset) \odot \emptyset \vdash t : A \\ & \text{SSF}((\Delta \mid \sigma_1 \mid \sigma_2) \odot \Gamma \vdash t : A) \implies \text{SSF}((\Delta, \mathbf{0} \mid \sigma_1, 0 \mid \sigma_2, r) \odot \Gamma, x : \mathbf{Type}_l \vdash t : A) \\ & \text{SSF}((\Delta \mid \sigma_1 \mid \sigma_2) \odot \Gamma \vdash t : A) \wedge \mathbf{Type}_l \not\leq^{+ve} B \\ & \implies \text{SSF}((\Delta, \sigma_3 \mid \sigma_1, s \mid \sigma_2, 0) \odot \Gamma, x : B \vdash t : A) \end{aligned}$$

By $\mathbf{Type}_l \not\leq^{+ve} B$ we mean \mathbf{Type}_l is not a positive subterm of B , avoiding higher-order typing terms (e.g., type constructors) which do not exist in SSF.

Under this restriction, we give a type-directed encoding mapping derivations of GRTT to SSF: given a GRTT derivation of judgment $(\Delta \mid \sigma_1 \mid \sigma_2) \odot \Gamma \vdash t : A$ we have that $\exists t_s$ (an SSF term) such that there is a derivation of judgment $[\Gamma] \vdash t_s : [A]_\tau$ in SSF where we interpret a subset of GRTT terms A as types:

$$\begin{aligned} \llbracket x \rrbracket_\tau &= x \\ \llbracket \mathbf{Type}_l \rrbracket_\tau &= \star_l \\ \llbracket (x :_{(0,r)} \mathbf{Type}_l) \rightarrow B \rrbracket_\tau &= \forall x : \star_l. \llbracket B \rrbracket_\tau \quad \text{where } \mathbf{Type}_l \not\leq^{+ve} B \\ \llbracket (x :_{(s,0)} A) \rightarrow B \rrbracket_\tau &= \llbracket A \rrbracket_\tau \rightarrow \llbracket B \rrbracket_\tau \quad \text{where } \mathbf{Type}_l \not\leq^{+ve} A, B \end{aligned}$$

Thus, dependent functions with \mathbf{Type} parameters that are computationally irrelevant (subject grade 0) map to \forall types, and dependent functions with parameters irrelevant in types (subject-type grade 0) map to regular function types.

We elide the full details but sketch key parts where functions and applications are translated inductively (where \mathbf{Ty}_l is shorthand for \mathbf{Type}_l):

$$\begin{aligned} \frac{(\Delta, \sigma_1 \mid \sigma_2, 0 \mid \sigma_3, r) \odot \Gamma, x : \mathbf{Ty}_l \vdash t : B}{\llbracket (\Delta \mid \sigma_2 \mid \sigma_1 + \sigma_3) \odot \Gamma \vdash \lambda x.t : (x :_{(0,r)} \mathbf{Ty}_l) \rightarrow B \rrbracket} &= \frac{\llbracket \Gamma \rrbracket, x : \star_l \vdash t_s : \llbracket B \rrbracket_\tau}{\llbracket \Gamma \rrbracket \vdash A(x : \star_l).t_s : \forall x : \star_l. \llbracket B \rrbracket_\tau} \\ \frac{(\Delta, \sigma_1 \mid \sigma_2, s \mid \sigma_3, 0) \odot \Gamma, x : A \vdash t : B}{\llbracket (\Delta \mid \sigma_2 \mid \sigma_1 + \sigma_3) \odot \Gamma \vdash \lambda x.t : (x :_{(s,0)} A) \rightarrow B \rrbracket} &= \frac{\llbracket \Gamma \rrbracket, x : \llbracket A \rrbracket_\tau \vdash t_s : \llbracket B \rrbracket_\tau}{\llbracket \Gamma \rrbracket \vdash \lambda(x : \llbracket A \rrbracket_\tau).t_s : \llbracket A \rrbracket_\tau \rightarrow \llbracket B \rrbracket_\tau} \\ \frac{(\Delta \mid \sigma_2 \mid \sigma_1 + \sigma_3) \odot \Gamma \vdash t_1 : (x :_{(0,r)} \mathbf{Ty}_l) \rightarrow B \quad (\Delta \mid \sigma_4 \mid \sigma_1) \odot \Gamma \vdash t_2 : \mathbf{Ty}_l}{\llbracket (\Delta \mid \sigma_2 \mid \sigma_3 + r * \sigma_4) \odot \Gamma \vdash t_1 t_2 : [t_2/x]B \rrbracket} &= \frac{\llbracket \Gamma \rrbracket \vdash t_s : \forall(x : \star_l). \llbracket B \rrbracket_\tau \quad \llbracket \Gamma \rrbracket \vdash T : \star_l}{\llbracket \Gamma \rrbracket \vdash t_s[T] : [T/x]\llbracket B \rrbracket_\tau} \\ \frac{(\Delta \mid \sigma_2 \mid \sigma_1 + \sigma_3) \odot \Gamma \vdash t_1 : (x :_{(s,0)} A) \rightarrow B \quad (\Delta \mid \sigma_4 \mid \sigma_1) \odot \Gamma \vdash t_2 : A}{\llbracket (\Delta \mid \sigma_2 + s * \sigma_4 \mid \sigma_3) \odot \Gamma \vdash t_1 t_2 : [t_2/x]B \rrbracket} &= \frac{\llbracket \Gamma \rrbracket \vdash t_s : \llbracket A \rrbracket_\tau \rightarrow \llbracket B \rrbracket_\tau \quad \llbracket \Gamma \rrbracket \vdash t'_s : \llbracket A \rrbracket_\tau}{\llbracket \Gamma \rrbracket \vdash t_s t'_s : [t'_s/x]\llbracket B \rrbracket_\tau} \end{aligned}$$

In the last case, note the presence of $[t'_s/x]\llbracket B \rrbracket_\tau$. Reasoning under the context of the encoding, this is proven equivalent to $\llbracket B \rrbracket_\tau$ since the subject type grade is 0 and therefore use of x in B is irrelevant.

Theorem 2 (Soundness and completeness of SSF embedding). *Given $\text{SSF}((\Delta \mid \sigma_1 \mid \sigma_2) \odot \Gamma \vdash t : A)$ and t_a in SSF where $\llbracket (\Delta \mid \sigma_1 \mid \sigma_2) \odot \Gamma \vdash t : A \rrbracket = \llbracket \Gamma \rrbracket \vdash t_s : \llbracket A \rrbracket_\tau$ then for CBN reduction $\rightsquigarrow^{\text{SSF}}$ in Stratified System F :*

$$\begin{aligned} (\text{soundness}) \quad \forall t'. t \rightsquigarrow t' &\implies \exists t'_s, t_s \rightsquigarrow^{\text{SSF}} t'_s \\ &\quad \wedge \llbracket (\Delta \mid \sigma_1 \mid \sigma_2) \odot \Gamma \vdash t' : A \rrbracket = \llbracket \Gamma \rrbracket \vdash t'_s : \llbracket A \rrbracket_\tau \\ (\text{completeness}) \quad \forall t'_s, t_s \rightsquigarrow^{\text{SSF}} t'_s &\implies \exists t'. t \rightsquigarrow t' \\ &\quad \wedge \llbracket (\Delta \mid \sigma_1 \mid \sigma_2) \odot \Gamma \vdash t' : A \rrbracket = \llbracket \Gamma \rrbracket \vdash t'_s : \llbracket A \rrbracket_\tau \end{aligned}$$

Thus, we can capture parametricity in GRTT via the judicious use of 0 grading (at either the type or computational level) for quantitative semirings. This embedding is not possible from QTT since QTT variables graded with 0 may be used arbitrarily in the types; the embedding here relies on GRTT's 0 type-grade capturing absence in types for quantitative semirings.

3.5 Graded Modal Types and Non-dependent Linear Types

GRTT can embed the reasoning present in other graded modal type theories (which often have a linear base), for example the explicit semiring-graded necessity modality found in coeffect calculi [10,23] and Granule [46]. We can recover the axioms of a graded necessity modality (usually modelled by an exponential graded comonad [23]). For example, in **Gerty** the following are well typed:

```
countit : (a : (.0, .2) Type) -> (z : (.1, .0) [.1] a) -> a
countit = \a z -> case z of [y] -> y
comult : (a : (.0, .2) Type) -> (z : (.1, .0) [.6] a) -> [.2] ([.3] a)
comult = \a z -> case z of [y] -> [[y]]
```

corresponding to $\varepsilon : \Box_1 A \rightarrow A$ and $\delta_{r,s} : \Box_{r*s} A \rightarrow \Box_r(\Box_s A)$: operations of graded necessity / graded comonads. Since we cannot use arbitrary terms for grades in the implementation, we have picked some particular grades here for `comult`. First-class grading is future work, discussed in Section 6.

Linear functions can be captured as $A \multimap B \triangleq (x :_{(1,r)} A) \rightarrow B$ for an exact usage semiring. It is straightforward to characterise a subset of GRTT programs that maps to the linear λ -calculus akin to the encodings above. Thus, GRTT provides a suitable basis for studying both linear and non-linear theories alike.

4 Metatheory

We now study GRTT’s metatheory. We first explain how substitution presents itself in the theory, and how type preservation follows from a relationship between equality and reduction. We then show admissibility of graded structural rules for contraction, exchange, and weakening, and strong normalization.

4.1 Substitution

We introducing substitution for well-formed contexts and then typing.

Lemma 4 (Substitution for well-formed contexts). *If the following hold:*

1. $(\Delta \mid \sigma_2 \mid \sigma_1) \odot \Gamma_1 \vdash t : A$ and 2. $(\Delta, \sigma_1, \Delta') \odot \Gamma_1, x : A, \Gamma_2 \vdash$

Then: $\Delta, (\Delta' \setminus |\Delta| + (\Delta' / |\Delta|) * \sigma_2) \odot \Gamma_1, [t/x]\Gamma_2 \vdash$

That is, given $\Gamma_1, x : A, \Gamma_2$ is well-formed, we can cut out x by substituting t for x in Γ_2 , accounting for the new usage in the context grade vectors. The usage of Γ_1 in t is given by σ_2 , and the usage in A by σ_1 . When substituting, Δ remains the same, as Γ_1 is unchanged. However, to account for the usage in $[t/x]\Gamma_2$, we have to form a new context grade vector $\Delta' \setminus |\Delta| + (\Delta' / |\Delta|) * \sigma_2$.

The operation $\Delta' \setminus |\Delta|$ (pronounced ‘discard’) removes grades corresponding to x , by removing the grade at index $|\Delta|$ from each grade vector in Δ' . Everything previously used in the typing of x in the context must now be distributed across $[t/x]\Gamma_2$, which is done by adding on $(\Delta' / |\Delta|) * \sigma_2$, which uses $\Delta' / |\Delta|$ (pronounced ‘choose’) to produce a vector of grades, which correspond to the grades cut out in $\Delta' \setminus |\Delta|$. The multiplication of $(\Delta' / |\Delta|) * \sigma_2$ produces a context grade vector by scaling σ_2 by each element of $(\Delta' / |\Delta|)$. When adding vectors, if the sizes of the vectors are different, then the shorter vector is right-padded with zeroes. Thus $\Delta' \setminus |\Delta| + (\Delta' / |\Delta|) * \sigma_2$ can be read as ‘ Δ' without the grades corresponding to x , plus the usage of t scaled by the prior usage of x ’.

For example, given typing $(((), (1) \mid 0, 1 \mid 1, 0) \odot a : \mathbf{Type}, y : a \vdash y : a$ and well-formed context $(((), (1), (1, 0), (0, 0, 2)) \odot a : \mathbf{Type}, y : a, x : a, z : t' \vdash$, where t' uses x twice, we can substitute y for x . Therefore, let $\Gamma_1 = a : \mathbf{Type}, y : a$ thus $|\Gamma_1| = 2$ and $\Gamma_2 = z : x$ and $\Delta' = ((0, 0, 2))$ and $\sigma_1 = 1, 0$ and $\sigma_2 = 0, 1$. Then the context grade of the substitution $[y/x]\Gamma_2$ is calculated as:

$$((0, 0, 2)) \setminus |\Gamma_1| = ((0, 0)) \quad (((0, 1, 2)) / |\Gamma_1|) * \sigma_2 = (2) * (0, 1) = ((0, 2))$$

Thus the resulting judgment is $((), (1), (0, 2)) \odot a : \text{Type}, y : a, z : [y/x]t' \vdash$.

Lemma 5 (Substitution for typing). *If the following premises hold:*

1. $(\Delta \mid \sigma_2 \mid \sigma_1) \odot \Gamma_1 \vdash t : A$
2. $(\Delta, \sigma_1, \Delta' \mid \sigma_3, s, \sigma_4 \mid \sigma_5, r, \sigma_6) \odot \Gamma_1, x : A, \Gamma_2 \vdash t' : B$
3. $|\sigma_3| = |\sigma_5| = |\Gamma_1|$

Then
$$\left(\begin{array}{c} \Delta, (\Delta' \setminus |\Delta| + (\Delta' / |\Delta|) * \sigma_2) \\ (\sigma_3 + s * \sigma_2), \sigma_4 \\ (\sigma_5 + r * \sigma_2), \sigma_6 \end{array} \right) \odot \Gamma_1, [t/x]\Gamma_2 \vdash [t/x]t' : [t/x]B.$$

As with substitution for well-formed contexts, we account for the replacement of x with t in Γ_2 by ‘cutting out’ x from the context grade vectors, and adding on the grades required to form t , scaled by the grades that described x ’s usage. We additionally must account for the altered subject and subject-type usage. We do this in a similar manner, by taking, for example, the usage of Γ_1 in the subject (σ_3), and adding on the grades required to form t , scaled by the grade with which x was previously used (s). Subject-type grades are calculated similarly.

4.2 Type Preservation

Lemma 6. *Reduction implies equality* If $(\Delta \mid \sigma_1 \mid \sigma_2) \odot \Gamma \vdash t_1 : A$ and $t_1 \rightsquigarrow t_2$, then $(\Delta \mid \sigma_1 \mid \sigma_2) \odot \Gamma \vdash t_1 = t_2 : A$.

Lemma 7. *Equality inversion* If $(\Delta \mid \sigma_1 \mid \sigma_2) \odot \Gamma \vdash t_1 = t_2 : A$, then $(\Delta \mid \sigma_1 \mid \sigma_2) \odot \Gamma \vdash t_1 : A$ and $(\Delta \mid \sigma_1 \mid \sigma_2) \odot \Gamma \vdash t_2 : A$.

Lemma 8. *Type preservation* If $(\Delta \mid \sigma_1 \mid \sigma_2) \odot \Gamma \vdash t : A$ and $t \rightsquigarrow t'$, then $(\Delta \mid \sigma_1 \mid \sigma_2) \odot \Gamma \vdash t' : A$.

Proof. By Lemma 6 we have $(\Delta \mid \sigma_1 \mid \sigma_2) \odot \Gamma \vdash t = t' : A$, and therefore by Lemma 7 we have $(\Delta \mid \sigma_1 \mid \sigma_2) \odot \Gamma \vdash t' : A$, as required.

4.3 Structural Rules

We now consider the structural rules of *contraction*, *exchange*, and *weakening*.

Lemma 9 (Contraction). *The following rule is admissible:*

$$\frac{\left(\begin{array}{c} \Delta_1, \sigma_1, (\sigma_1, 0), \Delta_2 \\ \sigma_2, s_1, s_2, \sigma_3 \\ \sigma_4, r_1, r_2, \sigma_5 \end{array} \right) \odot \Gamma_1, x : A, y : A, \Gamma_2 \vdash t : B \quad |\Delta_1| = |\sigma_2| = |\sigma_4| = |\Gamma_1|}{\left(\begin{array}{c} \Delta_1, \sigma_1, \text{contr}(|\Delta_1|; \Delta_2) \\ \sigma_2, (s_1 + s_2), \sigma_3 \\ \sigma_4, (r_1 + r_2), \sigma_5 \end{array} \right) \odot \Gamma_1, z : A, [z, z/x, y]\Gamma_2 \vdash [z, z/x, y]t : [z, z/x, y]B} \text{CONTR}$$

The operation $\text{contr}(\pi; \Delta)$ contracts the elements at index π and $\pi + 1$ for each vector in Δ by combining them with the semiring addition, defined $\text{contr}(\pi; \Delta) = \Delta \setminus (\pi + 1) + \Delta / (\pi + 1) * (\mathbf{0}^\pi, 1)$. Admissibility follows from the semiring addition, which serves to contract dependencies, being threaded throughout the rules.

Lemma 10 (Exchange). *The following rule is admissible:*

$$\frac{x \notin \text{FV}(B) \quad |\Delta_1| = |\sigma_3| = |\sigma_5| = |\Gamma_1| \quad \left(\begin{array}{c} \Delta_1, \sigma_1, (\sigma_2, 0), \Delta_2 \\ \sigma_3, s_1, s_2, \sigma_4 \\ \sigma_5, r_1, r_2, \sigma_6 \end{array} \right) \odot \Gamma_1, x : A, y : B, \Gamma_2 \vdash t : C}{\left(\begin{array}{c} \Delta_1, \sigma_2, (\sigma_1, 0), \text{exch}(|\Delta_1|; \Delta_2) \\ \sigma_3, s_2, s_1, \sigma_4 \\ \sigma_5, r_2, r_1, \sigma_6 \end{array} \right) \odot \Gamma_1, y : B, x : A, \Gamma_2 \vdash t : C} \text{Exc}$$

Notice that if you strip away the vector fragment and sizing premise, this is exactly the form of exchange we would expect in a dependent type theory: if x and y are assumptions in a context typing $t : C$, and the type of y does not depend upon x , then we can type $t : C$ when we swap the order of x and y .

The action on grade vectors is simple: we swap the grades associated with each of the variables. For the context grade vector however, we must do two things: first, we capture the formation of A with σ_1 , and the formation of B with $\sigma_1, 0$ (indicating x being used with grade 0 in B), then swap these around, cutting the final grade from $\sigma_2, 0$, and adding 0 to the end of σ_1 to ensure correct sizing. Next, the operation $\text{exch}(|\Delta_1|; \Delta_2)$ swaps the element at index $|\Delta_1|$ (i.e., that corresponding to usage of x) with the element at index $|\Delta_1| + 1$ (corresponding to y) for every vector in Δ_2 ; this exchange operation ensures that usage in the trailing context is reordered appropriately.

Lemma 11 (Weakening). *The following rule is admissible:*

$$\frac{\begin{array}{c} (\Delta_1, \Delta_2 \mid \sigma_1, \sigma'_1 \mid \sigma_2, \sigma'_2) \odot \Gamma_1, \Gamma_2 \vdash t : B \\ (\Delta_1 \mid \sigma_3 \mid \mathbf{0}) \odot \Gamma_1 \vdash A : \text{Type}_t \end{array} \quad |\sigma_1| = |\sigma_2| = |\Gamma_1|}{(\Delta_1, \sigma_3, \text{ins}(|\Delta_1|; 0; \Delta_2) \mid \sigma_1, 0, \sigma'_1 \mid \sigma_2, 0, \sigma'_2) \odot \Gamma_1, x : A, \Gamma_2 \vdash t : B} \text{WEAK}$$

Weakening introduces irrelevant assumptions to a context. We do this by capturing the usage in the formation of the assumption's type with σ_3 to preserve the well-formedness of the context. We then indicate irrelevance of the assumption by grading with 0 in appropriate places. The operation $\text{ins}(\pi; s; \Delta)$ inserts the element s at index π for each σ in Δ , such that all elements preceding index π (in σ) keep their positions, and every element at index π or greater (in σ) will be shifted one index later in the new vector. The 0 grades in the subject and subject-type grade vector positions correspond to the absence of the irrelevant assumption from the subject and subject's type.

4.4 Strong Normalization

We adapt Geuvers' strong normalization proof for the Calculus of Constructions (CC) [24] to a fragment of GRTT (called $\text{GRTT}^{\{0,1\}}$) restricted to two universe levels and without variables of type Type_1 . This results in a less expressive system than full GRTT when it comes to higher kinds, but this is orthogonal to the main idea here of grading. We briefly overview the strong normalization proof; details can be found in the extended version [44]. Note this strong normalization result is with respect to β -reduction only (our semantics does not include η -reduction).

We use the proof technique of saturated sets, based on the reducibility candidates of Girard [29]. While $\text{GRTT}^{\{0,1\}}$ has a collapsed syntax we use judgments to break typing up into stages. We use these sets to match on whether a term is a kind, type, constructor, or a function (we will refer to these as terms).

Definition 3. Typing can be broken up into the following stages:

$$\begin{aligned} \text{Kind} &:= \{A \mid \exists \Delta, \sigma_1, \Gamma. (\Delta \mid \sigma_1 \mid \mathbf{0}) \odot \Gamma \vdash A : \text{Type}_1\} \\ \text{Type} &:= \{A \mid \exists \Delta, \sigma_1, \Gamma. (\Delta \mid \sigma_1 \mid \mathbf{0}) \odot \Gamma \vdash A : \text{Type}_0\} \\ \text{Con} &:= \{t \mid \exists \Delta, \sigma_1, \sigma_2, \Gamma, A. (\Delta \mid \sigma_1 \mid \sigma_2) \odot \Gamma \vdash t : A \wedge (\Delta \mid \sigma_2 \mid \mathbf{0}) \odot \Gamma \vdash A : \text{Type}_1\} \\ \text{Term} &:= \{t \mid \exists \Delta, \sigma_1, \sigma_2, \Gamma, A. (\Delta \mid \sigma_1 \mid \sigma_2) \odot \Gamma \vdash t : A \wedge (\Delta \mid \sigma_2 \mid \mathbf{0}) \odot \Gamma \vdash A : \text{Type}_0\} \end{aligned}$$

Lemma 12 (Classification). *We have $\text{Kind} \cap \text{Type} = \emptyset$ and $\text{Con} \cap \text{Term} = \emptyset$.*

The classification lemma states that we can safely case split over kinds and types, or constructors and terms without fear of an overlap occurring.

Saturated sets are essentially collections of strongly normalizing terms that are closed under β -reduction. The intuition behind this proof is that every typable program ends up in some saturated set, and hence, is strongly normalizing.

Definition 4. [Base terms and saturated terms] Informally, the set of base terms \mathcal{B} is inductively defined from variables and Type_0 and Type_1 , and compound terms over base \mathcal{B} and strongly normalising terms SN .

A set of terms X is *saturated* if $X \subset \text{SN}$, $\mathcal{B} \subset X$, and if $\text{red}_k t \in X$ and $t \in \text{SN}$, then $t \in X$. Thus saturated sets are closed under strongly normalizing terms with a *key redex*, denoted $\text{red}_k t$, which are redexes or a redex at the head of an elimination form. SAT denotes the collection of saturated sets.

Lemma 13 (SN saturated). *All saturated sets are non-empty; SN is saturated.*

Since $\text{GRTT}^{\{0,1\}}$ allows computation in types as well as in types, we separate the interpretations for kinds and types, where the former is a set of the latter.

Definition 5. For $A \in \text{Kind}$, the kind interpretation, $\mathcal{K}[A]$, is defined:

$$\begin{aligned} \mathcal{K}[\text{Type}_0] &= \text{SAT} & \mathcal{K}[(x :_{(s,r)} A) \rightarrow B] &= \{f \mid f : \mathcal{K}[A] \rightarrow \mathcal{K}[B]\}, \text{ if } A, B \in \text{Kind} \\ \mathcal{K}[\square_s A] &= \mathcal{K}[A] & \mathcal{K}[(x :_{(s,r)} A) \rightarrow B] &= \mathcal{K}[A], \text{ if } A \in \text{Kind}, B \in \text{Type} \\ & & \mathcal{K}[(x :_{(s,r)} A) \rightarrow B] &= \mathcal{K}[B], \text{ if } A \in \text{Type}, B \in \text{Kind} \\ & & \mathcal{K}[(x :_r A) \otimes B] &= \mathcal{K}[A] \times \mathcal{K}[B], \text{ if } A, B \in \text{Kind} \\ & & \mathcal{K}[(x :_r A) \otimes B] &= \mathcal{K}[A], \text{ if } A \in \text{Kind}, B \in \text{Type} \\ & & \mathcal{K}[(x :_r A) \otimes B] &= \mathcal{K}[B], \text{ if } A \in \text{Type}, B \in \text{Kind} \end{aligned}$$

Next we define the interpretation of types, which requires the interpretation to be parametric on an interpretation of type variables called a type evaluation. This is necessary to make the interpretation well-founded (first realized by Girard [29]).

Definition 6. *Type valuations*, $\Delta \odot \Gamma \models \varepsilon$, are defined as follows:

$$\frac{}{\emptyset \odot \emptyset \models \emptyset} \text{E} \quad \frac{X \in \mathcal{K}[A] \quad \Delta \odot \Gamma \models \varepsilon}{(\Delta, \sigma \mid \mathbf{0}) \odot \Gamma \vdash A : \text{Type}_1} \text{TY} \quad \frac{\Delta \odot \Gamma \models \varepsilon}{(\Delta \mid \sigma \mid \mathbf{0}) \odot \Gamma \vdash A : \text{Type}_0} \text{TM}$$

Type valuations ignore term variables (rule TM), in fact, the interpretations of both types and kinds ignores them because we are defining sets of terms over types, and thus terms in types do not contribute to the definition of these sets. However as these interpretations define sets of open terms we must carry a graded context around where necessary. Thus, type valuations are with respect to a well-formed graded context $\Delta \odot \Gamma$. We now outline the type interpretation.

Definition 7. For type valuation $\Delta \odot \Gamma \models \varepsilon$ and a type $A \in (\text{Kind} \cup \text{Type} \cup \text{Con})$ with A typable in $\Delta \odot \Gamma$, the interpretation of types $\llbracket A \rrbracket_\varepsilon$ is defined inductively. For brevity, we list just a few illustrative cases, including modalities and some function cases; the complete definition is given in the extended version [44].

$$\begin{aligned}
 \llbracket \text{Type}_1 \rrbracket_\varepsilon &= \text{SN} \\
 \llbracket \text{Type}_0 \rrbracket_\varepsilon &= \lambda X \in \text{SAT.SN} \\
 \llbracket x \rrbracket_\varepsilon &= \varepsilon x && \text{if } x \in \text{Con} \\
 \llbracket \square_s A \rrbracket_\varepsilon &= \llbracket A \rrbracket_\varepsilon \\
 \llbracket \lambda x : A. B \rrbracket_\varepsilon &= \lambda X \in \mathcal{K}[\llbracket A \rrbracket_\varepsilon]. \llbracket B \rrbracket_{\varepsilon[x \mapsto X]} && \text{if } A \in \text{Kind}, B \in \text{Con} \\
 \llbracket A B \rrbracket_\varepsilon &= \llbracket A \rrbracket_\varepsilon (\llbracket B \rrbracket_\varepsilon) && \text{if } B \in \text{Con} \\
 \llbracket (x :_{(s,r)} A) \rightarrow B \rrbracket_\varepsilon &= \lambda X \in \mathcal{K}[\llbracket A \rrbracket_\varepsilon] \rightarrow \mathcal{K}[\llbracket B \rrbracket_\varepsilon]. \bigcap_{Y \in \mathcal{K}[\llbracket A \rrbracket_\varepsilon]} (\llbracket A \rrbracket_\varepsilon Y \rightarrow \llbracket B \rrbracket_{\varepsilon[x \mapsto Y]}(X(Y))) \\
 &&& \text{if } A, B \in \text{Kind}
 \end{aligned}$$

Grades play no role in the reduction relation for GRTT, and hence, our interpretation erases graded modalities and their introductory and elimination forms (translated into substitutions). In fact, the above interpretation can be seen as a translation of $\text{GRTT}^{\{0,1\}}$ into non-substructural set theory; there is no data-usage tracking in the image of the interpretation. Tensors are translated into Cartesian products whose eliminators are translated into substitutions similarly to graded modalities. All terms however remain well-typed through the interpretation.

The interpretation of terms corresponds to term valuations that are used to close the term before interpreting it into the interpretation of its type.

Definition 8. *Valid term valuations*, $\Delta \odot \Gamma \models_\varepsilon \rho$, are defined as follows:

$$\frac{}{\emptyset \odot \emptyset \models_\emptyset \emptyset} \text{E} \frac{\Delta \odot \Gamma \models_\varepsilon \rho \quad (\Delta \mid \sigma \mid \mathbf{0}) \odot \Gamma \vdash A : \text{Type}_1}{(\Delta, \sigma) \odot \Gamma, x : A \models_\varepsilon \rho[x \mapsto t]} \text{TY} \frac{t \in \llbracket A \rrbracket_\varepsilon \quad \Delta \odot \Gamma \models_\varepsilon \rho \quad (\Delta \mid \sigma \mid \mathbf{0}) \odot \Gamma \vdash A : \text{Type}_0}{(\Delta, \sigma) \odot \Gamma, x : A \models_\varepsilon \rho[x \mapsto t]} \text{TM}$$

We interpret terms as substitutions, but graded modalities must be erased and their elimination forms converted into substitutions (and similarly for the eliminator for tensor products).

Definition 9. Suppose $\Delta \odot \Gamma \models_\varepsilon \rho$. Then the *interpretation of a term* t typable in $\Delta \odot \Gamma$ is $\langle t \rangle_\rho = \rho t$, but where all let-expressions are translated into substitutions, and all graded modalities are erased.

Finally, we prove our main result using semantic typing which will imply strong normalization. Suppose $(\Delta \mid \sigma_1 \mid \sigma_2) \odot \Gamma \vdash t : A$, then:

Definition 10. *Semantic typing*, $(\Delta \mid \sigma_1 \mid \sigma_2) \odot \Gamma \models t : A$, is defined as follows:

1. If $(\Delta \mid \sigma \mid \mathbf{0}) \odot \Gamma \vdash A : \mathbf{Type}_1$, then for every $\Delta \odot \Gamma \models_\varepsilon \rho$, $(t)_\rho \in \llbracket A \rrbracket_\varepsilon$.
2. If $(\Delta \mid \sigma \mid \mathbf{0}) \odot \Gamma \vdash A : \mathbf{Type}_0$, then for every $\Delta \odot \Gamma \models_\varepsilon \rho$, $(t)_\rho \in \llbracket A \rrbracket_\varepsilon$.

Theorem 3 (Soundness for Semantic Typing). $(\Delta \mid \sigma_1 \mid \sigma_2) \odot \Gamma \models t : A$.

Corollary 1 (Strong Normalization). *We have $t \in \text{SN}$.*

5 Implementation

Our implementation **Gerty** is based on a bidirectionalised version of the typing rules here, somewhat following traditional schemes of bidirectional typing [19,20] but with grading (similar to Granule [46] but adapted considerably for the dependent setting). We briefly outline the implementation scheme and highlight a few key points, rules, and examples. We use this implementation to explore further applications of GRTT, namely optimising type checking algorithms.

Bidirectional typing splits declarative typing rules into *check* and *infer* modes. Furthermore, bidirectional GRTT rules split the grading context (left of \odot) into *input* and *output* contexts where $(\Delta \mid \sigma_1 \mid \sigma_2) \odot \Gamma \vdash t : A$ is implemented via:

$$(check) \Delta; \Gamma \vdash t \Leftarrow A; \sigma_1; \sigma_2 \quad or \quad (infer) \Delta; \Gamma \vdash t \Rightarrow A; \sigma_1; \sigma_2$$

where \Leftarrow rules *check* that t has type A and \Rightarrow rules *infer* (calculate) that t has type A . In both judgments, the context grading Δ and context Γ left of \vdash are inputs whereas the grade vectors σ_1 and σ_2 to the right of A are outputs. This input-output context approach resembles that employed in linear type checking [5,32,62]. Rather than following a “left over” scheme as in these works (where the output context explains what resources are left), the output grades here explain what has been used according to the analysis of grading (‘adding up’ rather than ‘taking away’).

For example, the following is the *infer* rule for function elimination:

$$\frac{\begin{array}{l} \Delta; \Gamma \vdash t_1 \Rightarrow (x :_{(s,r)} A) \rightarrow B; \sigma_2; \sigma_{13} \\ \Delta; \Gamma \vdash t_2 \Leftarrow A; \sigma_4; \sigma_1 \\ \Delta, \sigma_1; \Gamma, x : A \vdash B \Rightarrow \mathbf{Type}_l; \sigma_3, r; \mathbf{0} \quad \sigma_{13} = \sigma_1 + \sigma_3 \end{array}}{\Delta; \Gamma \vdash t_1 t_2 \Rightarrow [t_2/x]B; \sigma_2 + s * \sigma_4; \sigma_3 + r * \sigma_4} \Rightarrow \lambda_e$$

The rule can be read by starting at the input of the conclusion (left of \vdash), then reading top down through each premise, to calculate the output grades in the rule’s conclusion. Any concrete value or already-bound variable appearing in the output grades of a premise can be read as causing an equality check in the type checker. The last premise checks that the output subject-type grade σ_{13} from the first premise matches $\sigma_1 + \sigma_3$ (which were calculated by later premises).

In contrast, function introduction is a *check* rule:

$$\frac{\Delta; \Gamma \vdash A \Rightarrow \mathbf{Type}_l; \sigma_1; \mathbf{0} \quad \Delta, \sigma_1; \Gamma, x : A \vdash t \Leftarrow B; \sigma_2, s; \sigma_3, r}{\Delta; \Gamma \vdash \lambda x.t \Leftarrow (x :_{(s,r)} A) \rightarrow B; \sigma_2; \sigma_1 + \sigma_3} \Leftarrow \lambda_i$$

Thus, dependent functions can be checked against type $(x :_{(s,r)} A) \rightarrow B$ given input $\Delta; \Gamma$ by first inferring the type of A and checking that its output subject-type grade comprises all zeros $\mathbf{0}$. Then the body of the function t is checked against B under the context $\Delta, \sigma_1; \Gamma, x : A$ producing grade vectors σ_2, s' and σ_1, r' where it is checked that $s = s'$ and $r = r'$ (described implicitly in the rule), i.e., the calculated grades match those of the binder.

The implementation anticipates some further work for GRIT: the potential for grades which are first-class terms, for which we anticipate complex equations on grades. For grade equality, **Gerty** has two modes: one which normalises terms and then compares for syntactic equality, and the other which discharges constraints via an off-the-shelf SMT solver (we use Z3 [17]). We discuss briefly some performance implications in the next section.

Using Grades to Optimise Type Checking Abel posited that a dependent theory with quantitative resource tracking at the type level could leverage linearity-like optimisations in type checking [2]. Our implementation provides a research vehicle for exploring this idea; we consider one possible optimisation here.

Key to dependent type checking is the substitution of terms into types in elimination forms (i.e., application, tensor elimination). However, in a quantitative semiring setting, if a variable has 0 subject-type grade, then we know it is irrelevant to type formation (it is not semantically depended upon, i.e., during normalisation). Subsequently, substitutions into a 0-graded variable can be elided (or allocations to a closure environment can be avoided). We implemented this optimisation in **Gerty** when inferring the type of an application for $t_1 t_2$ (rule $\Rightarrow \lambda_e$ above), where the type of t_1 is inferred as $(x :_{(s,0)} A) \rightarrow B$. For a quantitative semiring we know that x irrelevant in B , thus we need not perform the substitution $[t_2/x]B$ when type checking the application.

We evaluate this on simple **Gerty** programs of an n -ary “fanout” combinator implemented via an n -ary application combinator, e.g., for arity 3:

```

app3 : (a : (0, 6) Type 0) -> (b : (0, 2) Type 0)
-> (x0 : (1, 0) a) -> (x1 : (1, 0) a) -> (x2 : (1, 0) a)
-> (f : (1, 0) ((y0 : (1,0) a) -> (y1 : (1,0) a) -> (y2 : (1,0) a) -> b)) -> b
app3 = \a -> \b -> \x0 -> \x1 -> \x2 -> \f -> f x0 x1 x2

fan3 : (a : (0, 4) Type 0) -> (b : (0, 2) Type 0)
-> (f : (1,0) ((z0 : (1,0) a) -> (z1 : (1,0) a) -> (z2 : (1,0) a) -> b))
-> (x : (3, 0) a) -> b
fan3 = \a -> \b -> \f -> \x -> app3 a b x x x f

```

Note that **fan3** uses its parameter x three times (hence the grade 3) which then incurs substitutions into the type of **app3** during type checking, but each such substitution is redundant since the type does not depend on these parameters, as reflected by the 0 subject-type grades.

To evaluate the optimisation and SMT solving vs. normalisation-based equality, we ran **Gerty** on the fan out program for arities from 3 to 8, with and without the optimisation and under the two equality approaches.

| n | Normalisation | | | SMT | | |
|-----|------------------|-----------------|---------|------------------|-----------------|---------|
| | Base ms | Optimised ms | Speedup | Base ms | Optimised ms | Speedup |
| 3 | 45.71 (1.72) | 44.08 (1.28) | 1.04 | 77.12 (2.65) | 76.91 (2.36) | 1.00 |
| 4 | 108.75 (4.09) | 89.73 (4.73) | 1.21 | 136.18 (5.23) | 162.95 (3.62) | 0.84 |
| 5 | 190.57 (8.31) | 191.25 (8.13) | 1.00 | 279.49 (15.73) | 289.73 (23.30) | 0.96 |
| 6 | 552.11 (29.00) | 445.26 (23.50) | 1.24 | 680.11 (16.28) | 557.08 (13.87) | 1.22 |
| 7 | 1821.49 (49.44) | 1348.85 (26.37) | 1.35 | 1797.09 (43.53) | 1368.45 (20.16) | 1.31 |
| 8 | 6059.30 (132.01) | 4403.10 (86.57) | 1.38 | 5913.06 (118.83) | 4396.90 (59.82) | 1.34 |

Table 1. Performance analysis of grade-based optimisations to type checking. Times in milliseconds to 2 d.p. with the standard error given in brackets. Measurements are the mean of 10 trials (run on a 2.7 Ghz Intel Core, 8Gb of RAM, Z3 4.8.8).

Table 1 gives the results. For grade equality by normalisation, the optimisation has a positive effect on speedup, getting increasingly significant (up to 38%) as the overall cost increases. For SMT-based grade equality, the optimisation causes some slow down for arity 4 and 5 (and just breaking even for arity 3). This is because working out whether the optimisation can be applied requires checking whether grades are equal to 0, which incurs extra SMT solver calls. Eventually, this cost is outweighed by the time saved by reducing substitutions. Since the grades here are all relatively simple, it is usually more efficient for the type checker to normalise and compare terms rather than compiling to SMT and starting up the external solver, as seen by longer times for the SMT approach.

The baseline performance here is poor (the implementation is not highly optimised) partly due to the overhead of computing type formation judgments often to accurately account for grading. However, such checks are often recomputed and could be optimised away by memoisation. Nevertheless this experiment gives the evidence that grades can indeed be used to optimise type checking. A thorough investigation of grade-directed optimisations is future work.

6 Discussion

Grading, Coeffects, and Quantitative Types The notion of *coeffects*, describing how a program depends on its context, arose in the literature from two directions: as a dualisation of effect types [48,49] and a generalisation of Bounded Linear Logic to general resource semirings [25,10]. Coeffect systems can capture reuse bounds, information flow security [23], hardware scheduling constraints [25], and sensitivity for differential privacy [16,22]. A coeffect-style approach also enables linear types to be retrofitted to Haskell [8]. A common thread is the annotation of variables in the context with usage information, drawn from a semiring. Our approach generalises this idea to capture type, context, and computational usage.

McBride [43] reconciles linear and dependent types, allowing types to depend on linear values, refined by Atkey [6] as Quantitative Type Theory. QTT employs coeffect-style annotation of each assumption in a context with an element of a resource accounting algebra, with judgments of the form:

$$x_1 \overset{\rho_1}{\vdash} A_1, \dots, x_n \overset{\rho_n}{\vdash} A_n \vdash M \overset{\rho}{\vdash} B$$

where ρ_i, ρ are elements of a semiring, and $\rho = 0$ or $\rho = 1$, respectively denoting a term which can be used in type formation (erased at runtime) or at runtime. Dependent function arrows are of the form $(x :^{\rho} A) \rightarrow B$, where ρ is a semiring element that denotes the computational usage of the parameter.

Variables used for type formation but not computation are annotated by 0. Subsequently, type formation rules are all of the form $0\Gamma \vdash T$, meaning every variable assumption has a 0 annotation. GRTT is similar to QTT, but differs in its more extensive grading to track usage in types, rather than blanketing all type usage with 0. In Atkey’s formulation, a term can be promoted to a type if its result and dependency quantities are all 0. A set of rules provide formation of computational type terms, but these are also graded at 0. Subsequently, it is not possible to construct an inhabitant of **Type** that can be used at runtime. We avoid this shortcoming allowing matching on types. For example, a computation t that inspects a type variable a would be typed as: $(\Delta, \mathbf{0}, \Delta' \mid \sigma_1, 1, \sigma'_1 \mid \sigma_2, r, \sigma'_2) \odot \Gamma, a : \mathbf{Type}, \Gamma' \vdash t : B$ denoting 1 computational use and r type uses in B .

At first glance, it seems QTT could be encoded into GRTT taking the semiring \mathcal{R} of QTT and parameterising GRTT by the semiring $\mathcal{R} \cup \{\hat{0}\}$ where $\hat{0}$ denotes arbitrary usage in type formation. However, there is impedance between the two systems as QTT always annotates type use with 0. It is not clear how to make this happen in GRTT whilst still having non-0 tracking at the computational level, since we use one semiring for both. Exploring an encoding is future work.

Choudhury et al. [13] give a system closely related (but arguably simpler) to QTT called GRAD. One key difference is that rather than annotating type usage with 0, grades are simply ignored in types. This makes for a surprisingly flexible system. In addition, they show that irrelevance is captured by the 0 grade using a heap-based semantics (a result leveraged in Section 3). GRAD however does not have the power of type-grades presented here.

Dependent Types and Modalities Dal Lago and Gaboardi extend PCF with linear and lightweight dependent types [15] (then adapted for differential privacy analysis [22]). They add a natural number type indexed by upper and lower bound terms which index a modality. Combined with linear arrows of the form $[a < I].\sigma \multimap \tau$ these describe functions using the parameter at most I times (where the modality acts as a binder for index variable a which denotes instantiations). Their system is leveraged to give fine-grained cost analyses in the context of Implicit Computational Complexity. Whilst a powerful system, their approach is restricted in terms of dependency, where only a specialised type can depend on specialised natural-number indexed terms (which are non-linear).

Gratzer et al. define a dependently-typed language with a Fitch-style modality [30]. It seems that such an approach could also be generalised to a graded modality, although we have used the natural-deduction style for our graded modality rather than the Fitch-style.

As discussed in Section 1, our approach closely resembles Abel’s *resourceful dependent types* [2]. Our work expands on the idea, including tensors and the graded modalities. We considerably developed the associated metatheory, provide an implementation, and study applications.

Further Work One expressive extension is to capture analyses which have an ordering, e.g., grading by a *pre-ordered* semiring, allowing a notion of *approximation*. This would enable analyses such as bounded reuse from Bounded Linear Logic [28], intervals with least- and upper-bounds on use [46], and top-completed semirings, with an ∞ -element denoting arbitrary usage as a fall-back. We have made progress into exploring the interaction between approximation and dependent types, and the remainder of this is left as future work.

A powerful extension of GRTT for future work is to allow grades to be first-class terms. Typing rules in GRTT involving grades could be adapted to internalise the elements as first-class terms. We could then, e.g., define the map function over sized vectors, which requires that the parameter function is used exactly the same number of times as the length of the vector:

$$\begin{aligned} \text{map} : (n :_{(0,5)} \text{nat}) \rightarrow (a :_{(0,n+1)} \text{Type}) \rightarrow (b :_{(0,n+1)} \text{Type}) \rightarrow \\ (f :_{(n,0)} (x :_{(1,0)} a) \rightarrow b) \rightarrow (xs :_{(1,0)} \text{Vec } n \ a) \rightarrow \text{Vec } n \ b \end{aligned}$$

This type provides strong guarantees: the only well-typed implementations do the correct thing, up to permutations of the result vector. Without the grading, an implementation could apply f fewer than n times, replicating some of the transformed elements; here we know that f must be applied exactly n -times.

A further appealing possibility for GRTT is to allow the semiring to be defined internally, rather than as a meta-level parameter, leveraging dependent types for proofs of key properties. An implementation could specify what is required for a semiring instance, e.g., a record type capturing the operations and properties of a semiring. The rules of GRTT could then be extended, similarly to the extension to first-class grades, with the provision of the semiring(s) coming from GRTT terms. Thus, anywhere with a grading premise $(\Delta \mid \sigma_1 \mid \sigma_2) \odot \Gamma \vdash r : \mathcal{R}$ would also require a premise $(\Delta \mid \sigma_2 \mid \mathbf{0}) \odot \Gamma \vdash \mathcal{R} : \text{Semiring}$. This opens up the ability for programmers and library developers to provide custom modes of resource tracking with their libraries, allowing domain-specific program verification.

Conclusions The paradigm of ‘grading’ exposes the inherent structure of a type theory, proof theory, or semantics by matching the underlying structure with some algebraic structure augmenting the types. This idea has been employed for reasoning about side effects via graded monads [35], and reasoning about data flow as discussed here by semiring grading. Richer algebras could be employed to capture other aspects, such as *ordered logics* in which the exchange rule can be controlled via grading (existing work has done this via modalities [34]).

We developed the core of grading in the context of dependent-types, treating types and terms equally (as one comes to expect in dependent-type theories). The tracking of data flow in types appears complex since we must account for how variables are used to form types in both the context and in the subject type, making sure not to repeat context formation use. The result however is a powerful system for studying dependencies in type theories, as shown by our ability to study different theories just by specialising grades. Whilst not yet a fully fledged implementation, **Gerty** is a useful test bed for further exploration.

Acknowledgments Orchard is supported by EPSRC grant EP/T013516/1.

References

1. Abadi, M., Banerjee, A., Heintze, N., Riecke, J.: A Core Calculus of Dependency. In: POPL. ACM (1999). <https://doi.org/10.1145/292540.292555>
2. Abel, A.: Resourceful Dependent Types. In: 24th International Conference on Types for Proofs and Programs, Abstracts (2018)
3. Abel, A., Bernardy, J.: A unified view of modalities in type systems. Proc. ACM Program. Lang. 4(ICFP), 90:1–90:28 (2020). <https://doi.org/10.1145/3408972>
4. Abel, A., Scherer, G.: On irrelevance and algorithmic equality in predicative type theory. Log. Methods Comput. Sci. 8(1) (2012). [https://doi.org/10.2168/LMCS-8\(1:29\)2012](https://doi.org/10.2168/LMCS-8(1:29)2012)
5. Allais, G.: Typing with Leftovers - A mechanization of Intuitionistic Multiplicative-Additive Linear Logic. In: Abel, A., Forsberg, F.N., Kaposi, A. (eds.) 23rd International Conference on Types for Proofs and Programs (TYPES 2017). Leibniz International Proceedings in Informatics (LIPIcs), vol. 104, pp. 1:1–1:22. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany (2018). <https://doi.org/10.4230/LIPIcs.TYPES.2017.1>
6. Atkey, R.: Syntax and Semantics of Quantitative Type Theory. In: Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2018, Oxford, UK, July 09–12, 2018. pp. 56–65 (2018). <https://doi.org/10.1145/3209108.3209189>
7. Atkey, R., Ghani, N., Johann, P.: A relationally parametric model of dependent type theory. In: Proceedings of the 41st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages. pp. 503–515 (2014). <https://doi.org/10.1145/2535838.2535852>
8. Bernardy, J.P., Boespflug, M., Newton, R.R., Peyton Jones, S., Spiwack, A.: Linear Haskell: practical linearity in a higher-order polymorphic language. Proceedings of the ACM on Programming Languages 2(POPL), 5 (2017). <https://doi.org/10.1145/3158093>
9. Brady, E., McBride, C., McKinna, J.: Inductive families need not store their indices. In: International Workshop on Types for Proofs and Programs. pp. 115–129. Springer (2003). https://doi.org/10.1007/978-3-540-24849-1_8
10. Brunel, A., Gaboardi, M., Mazza, D., Zdancewic, S.: A Core Quantitative Coeffect Calculus. In: Shao, Z. (ed.) Programming Languages and Systems - 23rd European Symposium on Programming, ESOP 2014, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2014, Grenoble, France, April 5–13, 2014, Proceedings. Lecture Notes in Computer Science, vol. 8410, pp. 351–370. Springer (2014). https://doi.org/10.1007/978-3-642-54833-8_19
11. Cardelli, L., Wegner, P.: On understanding types, data abstraction, and polymorphism. ACM Computing Surveys 17(4), 471–523 (Dec 1985). <https://doi.org/10.1145/6041.6042>
12. Cervesato, I., Pfenning, F.: A linear logical framework. Information and Computation 179(1), 19–75 (2002). <https://doi.org/10.1109/LICS.1996.561339>
13. Choudhury, P., Eades III, H., Eisenberg, R.A., Weirich, S.: A Graded Dependent Type System with a Usage-Aware Semantics. Proc. ACM Program. Lang. 5(POPL) (Jan 2021). <https://doi.org/10.1145/3434331>
14. Coquand, T., Huet, G.: The Calculus of Constructions. Ph.D. thesis, INRIA (1986)
15. Dal Lago, U., Gaboardi, M.: Linear dependent types and relative completeness. In: Logic in Computer Science (LICS), 2011 26th Annual IEEE Symposium on. pp. 133–142. IEEE (2011). <https://doi.org/10.1109/LICS.2011.22>

16. De Amorim, A.A., Gaboardi, M., Gallego Arias, E.J., Hsu, J.: Really Natural Linear Indexed Type Checking. In: Proceedings of the 26nd 2014 International Symposium on Implementation and Application of Functional Languages. p. 5. ACM (2014). <https://doi.org/10.1145/2746325.2746335>
17. De Moura, L., Bjørner, N.: Z3: An efficient SMT solver. In: International conference on Tools and Algorithms for the Construction and Analysis of Systems. pp. 337–340. Springer (2008). https://doi.org/10.1007/978-3-540-78800-3_24
18. De Vries, E., Plasmeijer, R., Abrahamson, D.M.: Uniqueness typing simplified. In: Symposium on Implementation and Application of Functional Languages. pp. 201–218. Springer (2007). https://doi.org/10.1007/978-3-540-85373-2_12
19. Dunfield, J., Krishnaswami, N.R.: Sound and complete bidirectional typechecking for higher-rank polymorphism with existentials and indexed types. *PACMPL* **3**(POPL), 9:1–9:28 (2019). <https://doi.org/10.1145/3290322>
20. Dunfield, J., Pfenning, F.: Tridirectional typechecking. In: Proceedings of the 31st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages. pp. 281–292. POPL '04, ACM, New York, NY, USA (2004). <https://doi.org/10.1145/964001.964025>
21. Eades, H., Stump, A.: Hereditary substitution for Stratified System F. In: International Workshop on Proof-Search in Type Theories, PSTT. vol. 10 (2010)
22. Gaboardi, M., Haeberlen, A., Hsu, J., Narayan, A., Pierce, B.C.: Linear dependent types for differential privacy. In: POPL. pp. 357–370 (2013). <https://doi.org/10.1145/2429069.2429113>
23. Gaboardi, M., Katsumata, S.y., Orchard, D., Breuvar, F., Uustalu, T.: Combining Effects and Coeffects via Grading. In: Proceedings of the 21st ACM SIGPLAN International Conference on Functional Programming. p. 476–489. ICFP 2016, Association for Computing Machinery, New York, NY, USA (2016). <https://doi.org/10.1145/2951913.2951939>
24. Geuvers, H.: A short and flexible proof of strong normalization for the calculus of constructions. In: Dybjer, P., Nordström, B., Smith, J. (eds.) *Types for Proofs and Programs*. pp. 14–38. Springer Berlin Heidelberg, Berlin, Heidelberg (1995). https://doi.org/10.1007/3-540-60579-7_2
25. Ghica, D.R., Smith, A.I.: Bounded linear types in a resource semiring. In: ESOP. pp. 331–350 (2014). https://doi.org/10.1007/978-3-642-54833-8_18
26. Girard, J.Y.: Une extension de l'interprétation de gödel a l'analyse, et son application a l'élimination des coupures dans l'analyse et la théorie des types. In: *Studies in Logic and the Foundations of Mathematics*, vol. 63, pp. 63–92. Elsevier (1971)
27. Girard, J.Y.: Linear logic. *Theoretical Computer Science* **50**, 1–102 (1987)
28. Girard, J.Y., Scedrov, A., Scott, P.J.: Bounded linear logic: a modular approach to polynomial-time computability. *Theoretical computer science* **97**(1), 1–66 (1992). [https://doi.org/10.1016/0304-3975\(92\)90386-T](https://doi.org/10.1016/0304-3975(92)90386-T)
29. Girard, J., Taylor, P., Lafont, Y.: *Proofs and types*, vol. 7. Cambridge University Press Cambridge (1989)
30. Gratzel, D., Sterling, J., Birkedal, L.: Implementing a modal dependent type theory. *Proc. ACM Program. Lang.* **3**(ICFP), 107:1–107:29 (2019). <https://doi.org/10.1145/3341711>
31. Henglein, F., Makhholm, H., Niss, H.: Effect types and region-based memory management. *Advanced Topics in Types and Programming Languages* pp. 87–135 (2005)
32. Hodas, J.S.: *Logic programming in intuitionistic linear logic: Theory, design and implementation*. PhD Thesis, University of Pennsylvania, Department of Computer and Information Science (1994)

33. Hofmann, M.: Syntax and semantics of dependent types. *Semantics and logics of computation* **14**, 79 (1997)
34. Jiang, J., Eades III, H., de Paiva, V.: On the lambek calculus with an exchange modality. In: Ehrhard, T., Fernández, M., de Paiva, V., de Falco, L.T. (eds.) *Proceedings Joint International Workshop on Linearity & Trends in Linear Logic and Applications, Linearity-TLLA@FLoC 2018*, Oxford, UK, 7-8 July 2018. *EPTCS*, vol. 292, pp. 43–89 (2018). <https://doi.org/10.4204/EPTCS.292.4>
35. Katsumata, S.: Parametric effect monads and semantics of effect systems. In: *Proceedings of POPL*. pp. 633–646. ACM (2014). <https://doi.org/10.1145/2535838.2535846>
36. Krishnaswami, N.R., Dreyer, D.: Internalizing relational parametricity in the extensional calculus of constructions. In: *Computer Science Logic 2013 (CSL 2013)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik (2013). <https://doi.org/10.4230/LIPIcs.CSL.2013.432>
37. Krishnaswami, N.R., Pradic, P., Benton, N.: Integrating linear and dependent types. In: *ACM SIGPLAN Notices*. vol. 50, pp. 17–30. ACM (2015)
38. Leivant, D.: Finitely stratified polymorphism. *Information and Computation* **93**(1), 93–113 (1991). [https://doi.org/10.1016/0890-5401\(91\)90053-5](https://doi.org/10.1016/0890-5401(91)90053-5)
39. Luo, Z., Zhang, Y.: A linear dependent type theory. *Types for Proofs and Programs (TYPES 2016)*, Novi Sad (2016)
40. Martin-Löf, P.: An Intuitionistic Theory of Types: Predicative Part. In: Rose, H.E., Shepherdson, J.C. (eds.) *Studies in Logic and the Foundations of Mathematics, Logic Colloquium '73*, vol. 80, pp. 73–118. Elsevier (Jan 1975). [https://doi.org/10.1016/S0049-237X\(08\)71945-1](https://doi.org/10.1016/S0049-237X(08)71945-1)
41. Martin-Löf, P.: *Intuitionistic Type Theory* (Jun 1980)
42. Martin-Löf, P.: Constructive Mathematics and Computer Programming. In: Cohen, L.J., Łoś, J., Pfeiffer, H., Podewski, K.P. (eds.) *Studies in Logic and the Foundations of Mathematics, Logic, Methodology and Philosophy of Science VI*, vol. 104, pp. 153–175. Elsevier (Jan 1982). [https://doi.org/10.1016/S0049-237X\(09\)70189-2](https://doi.org/10.1016/S0049-237X(09)70189-2)
43. McBride, C.: I Got Plenty o' Nuttin', pp. 207–233. Springer International Publishing, Cham (2016). https://doi.org/10.1007/978-3-319-30936-1_12
44. Moon, B., Eades III, H., Orchard, D.: Graded modal dependent type theory. *CoRR abs/2010.13163* (2020), <https://arxiv.org/abs/2010.13163>
45. Nuyts, A., Vezzosi, A., Devriese, D.: Parametric quantifiers for dependent type theory. *Proceedings of the ACM on Programming Languages* **1**(ICFP), 32:1–32:29 (Aug 2017). <https://doi.org/10.1145/3110276>
46. Orchard, D., Liepelt, V.B., Eades III, H.: Quantitative Program Reasoning with Graded Modal Types. *Proc. ACM Program. Lang.* **3**(ICFP), 110:1–110:30 (Jul 2019). <https://doi.org/10.1145/3341714>
47. Palmgren, E.: On universes in type theory. *Twenty-five years of constructive type theory* **36**, 191–204 (1998)
48. Petricek, T., Orchard, D., Mycroft, A.: Coeffects: Unified Static Analysis of Context-Dependence. In: *ICALP* (2). pp. 385–397 (2013). https://doi.org/10.1007/978-3-642-39212-2_35
49. Petricek, T., Orchard, D., Mycroft, A.: Coeffects: A calculus of context-dependent computation. In: *Proceedings of the 19th ACM SIGPLAN International Conference on Functional Programming*. pp. 123–135. *ICFP '14*, ACM (2014). <https://doi.org/10.1145/2628136.2628160>

50. Pfenning, F.: Intensionality, extensionality, and proof irrelevance in modal type theory. In: Proceedings 16th Annual IEEE Symposium on Logic in Computer Science. pp. 221–230. IEEE (2001). <https://doi.org/10.1109/LICS.2001.932499>
51. Reed, J.: Extending higher-order unification to support proof irrelevance. In: International Conference on Theorem Proving in Higher Order Logics. pp. 238–252. Springer (2003). https://doi.org/10.1007/10930755_16
52. Reynolds, J.C.: Towards a theory of type structure. In: Programming Symposium. pp. 408–425. Springer (1974). https://doi.org/10.1007/3-540-06859-7_148
53. Reynolds, J.C.: Types, abstraction and parametric polymorphism. In: Information Processing 83, Proceedings of the IFIP 9th World Computer Congress. pp. 513–523 (1983)
54. Shi, R., Xi, H.: A linear type system for multicore programming in ATS. *Science of Computer Programming* **78**(8), 1176–1192 (2013). <https://doi.org/10.1016/j.scico.2012.09.005>
55. Tofte, M., Talpin, J.P.: Region-based memory management. *Information and computation* **132**(2), 109–176 (1997). <https://doi.org/10.1006/inco.1996.2613>
56. Tov, J.A., Pucella, R.: Practical affine types. In: Proceedings of the 38th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2011, Austin, TX, USA, January 26–28, 2011. pp. 447–458 (2011). <https://doi.org/10.1145/1926385.1926436>
57. Wadler, P.: Theorems for free! In: Proceedings of the fourth international conference on Functional programming languages and computer architecture. pp. 347–359 (1989). <https://doi.org/10.1145/99370.99404>
58. Wadler, P.: Linear Types Can Change the World! In: *Programming Concepts and Methods*. North (1990)
59. Wadler, P.: Propositions as types. *Communications of the ACM* **58**(12), 75–84 (2015). <https://doi.org/10.1145/2699407>
60. Walker, D.: Substructural type systems. *Advanced Topics in Types and Programming Languages* pp. 3–44 (2005)
61. Wood, J., Atkey, R.: A Linear Algebra Approach to Linear Metatheory. *CoRR* **abs/2005.02247** (2020), <https://arxiv.org/abs/2005.02247>
62. Zalakain, U., Dardha, O.: Pi with leftovers: a mechanisation in Agda. *arXiv preprint arXiv:2005.05902* (2020)

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

