

The Flow of ODEs: Formalization of Variational Equation and Poincaré Map

Fabian Immler¹  · Christoph Traut¹

Received: 20 March 2017 / Accepted: 4 January 2018 / Published online: 23 January 2018
© Springer Science+Business Media B.V., part of Springer Nature 2018

Abstract Formal analysis of ordinary differential equations (ODEs) and dynamical systems requires a solid formalization of the underlying theory. The formalization needs to be at the correct level of abstraction, in order to avoid drowning in tedious reasoning about technical details. The *flow* of an ODE, i.e., the solution depending on initial conditions, and a dedicated type of bounded linear functions yield suitable abstractions. The dedicated type integrates well with the type-class based analysis in Isabelle/HOL and we prove advanced properties of the flow, most notably, differentiable dependence on initial conditions via the variational equation. Moreover, we formalize the notion of first return or Poincaré map and prove its differentiability. We provide rigorous numerical algorithm to solve the variational equation and compute the Poincaré map.

Keywords Isabelle/HOL · Analysis · Ordinary differential equation · Dynamical system · Poincaré map

1 Introduction

Ordinary differential equations (ODEs) are ubiquitous for modeling continuous problems in e.g., physics, biology, or economics. A formalization of the theory of ODEs allows us to verify algorithms for the analysis of such systems. A popular example, where a verified algorithm is highly relevant, is Tucker’s proof on the topic of a strange attractor for the Lorenz equations [26]. This proof relies on the output of a computer program, that computes bounds for analytical properties of the so-called Poincaré map of the *flow* of an ODE.

✉ Fabian Immler
immler@in.tum.de

Christoph Traut
trautc@in.tum.de

¹ Institut für Informatik, Technische Universität München, Garching, Germany

The flow is the solution as a function depending on an initial condition. We formalize the flow and prove conditions for analytical properties like continuity of differentiability (the derivative is of particular importance in Tucker's proof). Most of these properties seem very "natural", as Hirsch, Smale and Devaney call them in their textbook [11]. However, despite being "natural" properties and fairly standard results, they are delicate to prove: In the textbook, the authors present these properties rather early, but

"postpone all of the technicalities [. . .], primarily because understanding this material demands a firm and extensive background in the principles of real analysis".

In this article, we show that it is feasible to cope with these technicalities in a formal setting and confirm that Isabelle/HOL supplies a sufficient background of real analysis.

We present our Isabelle/HOL library for reasoning about the flow of ODEs. The main results are formalizations of continuous and differentiable dependence on initial conditions. The differentiable dependence is characterized by a particular ODE, the variational equation. The Poincaré map is a useful tool for reasoning about dynamical systems and is defined in terms of the flow. Usually, textbooks give a rigorous treatment of the Poincaré map only in a particular situation, namely in a local neighborhood of a periodic point. Tucker's proof, for example, requires a more general notion, for which we give a precise formalization here. We show how to use existing rigorous numerical algorithms to solve the variational equation and compute derivatives of the Poincaré map. The variational equation is posed on the space of linear functions. We introduce a separate type for this space in order to profit from the type class based formalization of mathematics in Isabelle/HOL.

We are not aware of any other formalization that covers this foundational part of the theory of ODEs in similar detail.

The remainder of this article is structured as follows: We will first (in Sect. 2) present the "interface" to our theory, i.e., the definitions and assumptions that are needed for formalizing our main results. Any potential user of the library needs in principle only know about these concepts. Section 3 follows with a high-level treatment of our formalization of the Poincaré map. Because the general topic is very theoretical and foundational work, we present a practical application with rigorous numerical computations right afterwards in Sect. 4.

Only then, we go into the details of the techniques that we used to make this formalization possible. Mathematics and analysis are formalized in Isabelle mostly based on type classes and filters, as has been presented earlier in earlier work [12]. We follow this path to formalize the foundations of our work:

Several proofs needed the notion of a uniform limit. We cast this notion into the "Isabelle/HOL approach to limits": we define it using a filter. This gives a versatile formalization and one can profit from the existing infrastructure for filters in limits. This will be presented in Sect. 5.

The derivative of the flow is a linear function. The space of linear functions forms a Banach space. In order to profit from the structure and properties that hold in a Banach space (which is a type class in Isabelle/HOL), we needed to introduce a type of bounded linear functions. We will present this type and further applications of its formalization in Sect. 6.

In Sect. 7, we present the technical lemmas that are needed to prove continuity and differentiability of the flow in order to give an impression of the kind of reasoning that is required. Section 8 contains a similarly detailed discussion about the Poincaré map.

All of the theorems we present here are formalized in Isabelle/HOL [23], the source code is available for Isabelle2017 in the Archive of Formal Proof [18].

This article is based on an earlier version that was presented at the conference ITP 2016 in Nancy, France [19]. It extends the conference paper with a treatment of the notion of Poincaré

map (Sects. 3, 8, and parts of Sect. 4) together with some necessary background on inverse and implicit functions (Sects. 6.5 and 6.6).

2 The Flow of a Differential Equation

In this section, we introduce the concept of *flow* and *existence interval* (which guarantees that the flow is well-defined) and present our main results (without proofs at first, we will present some of the lemmas leading to the proofs in Sect. 7).

The claim we want to make in this section is the flow as definition is a suitable abstraction for initial value problems. But beware and do not get deceived by simplicity of statements: as already mentioned in the introduction, these are all “natural” properties, but the proofs (also in the textbook) require many technical lemmas.

First of all, let us introduce the concepts we are interested in. We consider open sets T, X and an autonomous¹ ODE with right hand side f

$$\dot{x}(t) = f(x(t)), \text{ where } f : \mathbb{R}^n \rightarrow \mathbb{R}^n \text{ is a function from } X \text{ to } X \tag{1}$$

Under mild assumptions (which we will make more precise later in Definition 18), there exists a solution $\phi(t)$, which is unique for an initial condition $x(t_0) = x_0$. To emphasize the dependence on the initial condition, we write $\phi(x_0, t)$ for the solution of Eq. (1). This solution depending on initial conditions is called the *flow* of the differential equation.

Definition 1 (Flow) The *flow* $\phi(x_0, t)$ is the (unique) solution of the ODE (1) with initial condition $x(0) = x_0$

The solution does not necessarily exist for every $t \in T$. For example, solutions can *explode* in finite time s : if $\lim_{t \rightarrow s} \phi(z_0, t) = \infty$, then the flow is only defined for $t < s$ as is illustrated in Fig. 1 for $\phi(z_0, _)$. We therefore need to define a notion of (maximal) existence interval.

Definition 2 (Maximal Existence Interval) The *maximal existence interval* of the ODE (1) is the open interval

$$ex-ivl(x_0) :=]\alpha; \beta[$$

for $\alpha, \beta \in \mathbb{R} \cup \{\infty, -\infty\}$, such that $\phi(x_0, t)$ is a solution for $t \in ex-ivl(x_0)$. Moreover for every other interval I and every solution $\psi(x_0, t)$ for $t \in I$, one has $I \subseteq ex-ivl(x_0)$ and $\forall t \in I. \psi(x_0, t) = \phi(x_0, t)$.

We claim that the flow ϕ (together with *ex-ivl*, which guarantees the flow to be well-defined) is a convenient way to talk about solutions in a theorem prover. After guaranteeing that they are well-defined, these constants have nicely algebraic properties, which can be stated without further assumptions.

2.1 Composition of Solutions

The first nicely algebraic property is the abstract property of the generic notion of flow. This notion makes it possible to easily state composition of solutions and to algebraically reason

¹ This means that f does not depend on t . This restriction makes the presentation clearer. Many of our results are also formalized for non-autonomous ODEs. Moreover, for a large class of ODEs, a reduction to the autonomous case is possible.




Fig. 1 Illustration of some flow ϕ for initial values x_0, y_0, z_0 and times $s, s + t$

about them. As illustrated in Fig. 1, flowing from x_0 for time $s + t$ is equivalent to first flowing for time s , and from there flowing for time t .

This only works if the flow is defined also for the intermediate times (the theorem can not be true for $\phi(x_0, t + (-t))$ if $t \notin \text{ex-ivl}$).

Theorem 1 (Flow property)

$$s \in \text{ex-ivl}(x_0) \implies t \in \text{ex-ivl}(\phi(x_0, s)) \implies \phi(x_0, s + t) = \phi(\phi(x_0, s), t)$$

2.2 Continuity of the Flow

In the previous lemma, the assumption that the flow is defined (i.e., that the time is contained in the existence interval) was important. Let us now study the domain $\Omega = \{(x, t) \mid t \in \text{ex-ivl}(x)\} \subseteq X \times T$ of the flow in more detail. Ω is called the *state space*.

For the first “natural” property, we consider an element in the state space. $(x, t) \in \Omega$ means that we can follow a solution starting at x for time t . It is “natural” to expect that

solutions starting close to x can be followed for times that are close to t . In topological parlance, the state space is open.

Theorem 2 (Open State Space) *open Ω*

In the previous theorem, the state space allows us to reason about the fact that solutions are *defined* for close times and initial values. For quantifying how deviations in the initial values are propagated by the flow, Grönwall’s lemma is an important tool that is used in several proofs. Because of its importance in the theory of dynamical systems, we list it here as well, despite it being a rather technical result. Starting from an *implicit* inequality $g(t) \leq C + K \int_0^t g(s) ds$ involving a continuous, nonnegative function $g : \mathbb{R} \rightarrow \mathbb{R}$, it allows one to deduce an *explicit* bound for g .

Lemma 1 (Grönwall)

$$0 < C \implies 0 < K \implies \text{continuous-on } [0; a] \ g \implies$$

$$\forall t \in [0; a]. 0 \leq g(t) \leq C + K \int_0^t g(s) ds \implies$$

$$\forall t \in [0; a]. g(t) \leq C e^{Kt}$$

Grönwall’s lemma can be used to show that solutions deviate *at most* exponentially fast: $\exists K. |\phi(x, t) - \phi(y, t)| < |x - y|e^{K|t|}$ (see also Lemma 11). Therefore, by choosing x and y close enough, one can make the distance of the solutions arbitrarily small. This implies that the flow is continuous in its first argument and can be used to show that the flow is continuous on the state space.

Theorem 3 (Continuity of Flow) *continuous-on $\Omega \ \phi$*

2.3 Differentiability of the Flow

Continuity states that small deviation in the initial values result in small deviations of the flow. But one can be more precise on how initial deviations propagate. A nice property of the flow is that it is differentiable: the way initial deviations propagate can be approximated by a linear function. So instead of solving the ODE for perturbed initial values, one can approximate the resulting perturbation with the linear function: We write $\phi_t := \lambda x. \phi(x, t)$ and application of a linear function L to some argument x application as $L \cdot x$. Then one has $D\phi_t|_x \cdot v \approx \phi_t(x + v) - \phi_t(x)$. By using a basis vector for v , one gets the corresponding partial derivative of the flow. $D\phi_t|_x$ is called the space derivative of the flow.

As an example, Fig. 2 depicts a two-dimensional flow ϕ starting at (x_0, y_0) and its evolution (in black) up to time $t = 2$ in black. Along with the flow, it shows the evolution of the partial derivatives $\frac{\partial \phi((x_0, y_0), t)}{\partial x} = D\phi_t|(x_0, y_0) \cdot (1, 0)$ and $\frac{\partial \phi((x_0, y_0), t)}{\partial y} = D\phi_t|(x_0, y_0) \cdot (0, 1)$.

The derivative of ϕ with respect to time is given by the ODE f , one can therefore write the total derivative of the flow as $D\phi|(x, t) \cdot (d_x, d_t) := D\phi_t|_x \cdot d_x + d_t f(\phi(x, t))$. Our main result is the formalization of the fact that the total derivative of the flow exists and is continuous.

Theorem 4 (Differentiability of the Flow) *If f is continuously differentiable, then for every $(x, t) \in \Omega$ there exists a linear function $W(x, t)$, which is the derivative of the flow at (x, t) . Moreover W is continuous on the state space.*

$$\exists W. D\phi|(x, t) = W(x, t) \wedge \text{continuous-on } \Omega \ W$$




Fig. 2 Flow ϕ of the van der Pol system $(\dot{x}, \dot{y}) = (y, (1 - x^2)y - x)$ and its partial derivatives $\frac{\partial\phi}{\partial x}, \frac{\partial\phi}{\partial y}$ (scaled by the factor $\frac{1}{2}$) for initial condition $(x_0, y_0) = (1.4, 2.25)$

3 The Poincaré Map

The Poincaré map (or *first return map*) is an important tool for studying dynamical systems. The flow describes the evolution of a continuous system with respect to *time*. But often, one is not interested in the evolution with respect to *time*, but rather to some *space* variables and the Poincaré map allows to reason about this. To this end, one considers a continuous system with flow $\phi(x, t)$ and a Poincaré section. A Poincaré section is a subset Σ of the state space, which is usually given as an implicit surface $\Sigma = \{x \mid s(x) = c\}$ with continuously differentiable s .

The Poincaré map $P(x)$ is defined as the point where the flow starting from x first hits the Poincaré section Σ . To be more precise, let us consider the notion of first return time $\tau(x)$.

Definition 3 (Return Time) $\tau(x)$ is the least $t > 0$ such that $\phi(x, t) \in \Sigma$.

Obviously, τ is only well-defined for values that actually return to Σ , which we encode in the predicate *returns-to*.

Definition 4

$$\text{returns-to}(\Sigma, x) := \exists t > 0. \phi(x, t) \in \Sigma$$

The return time can then be used to define the Poincaré map as follows.

Definition 5 (*Poincaré map*)

$$P(x) := \phi(x, \tau(x))$$

It is interesting to note that our Definition of Poincaré map slightly differs from the approach that many textbooks (e.g., [24, chapter 3.4], [25, chapter 5.8], [11, chapter 10.3]) take. The original application of Poincaré maps is the study of periodic orbits, and textbooks usually define the return time implicitly for a periodic point as follows: For a periodic point x with period $\tau(x)$ and a choice of Σ transversal to the flow, these textbooks invoke the implicit function theorem to obtain a continuous function, which is declared to be the return time. This way, the definition implicitly depends on x and is valid only locally.

In contrast, our construction in Definition 5 (with the first return time τ according to Definition 3) yields a notion of Poincaré map that is given globally (on the whole state space and not just on Σ) and a-priori (without any implicit constructions): $P(x)$ is well-defined for all values x with *returns-to* (Σ, x) . In particular, x need not be element of Σ and also not in some (implicitly defined) neighborhood of a periodic point. P is well-defined even when there is no periodic point. Real applications like e.g., Tucker's proof require such a more flexible notion of Poincaré map.

In Sect. 8, we precise the assumptions under which τ (and therefore P) is continuous and differentiable.

4 Rigorous Numerics

In this section, we show that the formalization is not just something abstract, but rather something that we can use to specify the result of concrete computations: The space derivative $D\phi_t|_x$ of the flow can be characterized as the solution of a linear, matrix-valued ODE, a byproduct of the (constructive) proof of differentiability in Lemma 15: The space derivative $D\phi_t|_x$ is the solution W to the following ODE

$$\dot{W}(t) = Df|_{\phi(x_0,t)} \cdot W(t)$$

with the identity matrix I as initial condition $W(0) = I$.

We encode this matrix valued variational equation into a vector valued one, use an existing rigorous numerical algorithm for solving ODEs in Isabelle [14] to compute bounds on the solutions. Re-interpreting the result as bounds on matrices, we obtain bounds on the solution of the variational equation. As a concrete example, we use the van der Pol system: $\dot{x} = y$; $\dot{y} = (1 - x^2)y - x$ for the initial condition $x_0 = 1.4$ and $y_0 = 2.25$.

The overall setup for the computation is as follows: We have an executable specification of a second-order Runge–Kutta method (which is formally verified to produce rigorous enclosures for the solution of an ODE) and use Isabelle's code generator [8] to generate SML code from this specification. We choose to compute the evolution until time $t = 2$ with an adaptive stepsize controlling an absolute and relative error of about 2^{-12} . The computation takes about 15 s on an average laptop computer. Figure 2 was generated using the output of the verified algorithm. The corresponding theorem for the final inclusion is as follows (the theorem statement needs to be given by the user, it is proved by computing bounds using the aforementioned Runge–Kutta method, trusting code generation in SML):

Theorem 5

$$\begin{aligned} \phi((1.4, 2.25), 2) &\in ([0.1510; 0.1524], [-1.0353; -1.0334]) \\ D\phi_{2|_{(1.4, 2.25)}} &\in \begin{pmatrix} [0.2141; 0.2173] & [0.4262; 0.4276] \\ [-0.108; -0.110] & [0.2967; 0.2987] \end{pmatrix} \end{aligned}$$

In addition to the variational equation, we can also compute the Poincaré map P . In Fig. 3, we start with an initial set $[1.35; 1.45] \times \{2.25\}$ and aim to compute its Poincaré map returning to $y = 2.25$ and $x \in [1; 2]$. An intermediate Poincaré section at $x = -1$ helps in performing the computations: When the initial set evolves until about $x = -1$, one can see (in the magnified lower part of the figure) that the size of the reachable sets right of $x = -1$ is relatively large in the x -direction. The Poincaré section reduces the size in this direction to zero. Since the overapproximation error grows with the size of the reachable sets, such an operation helps to maintain precision and performance of the reachability analysis algorithm. The intersection is computed geometrically [15]. The framework that combines continuous reachability with such intersections is described in an earlier paper [16]. We use affine arithmetic to evaluate the expression for the derivative of the Poincaré map from theorem 13, where the discretization step size gives an estimate for the return time τ .

As a result, Fig. 3 was created using the output of the verified algorithm and this algorithm yields an enclosure for the Poincaré map P and its derivative DP :

Theorem 6

$$\begin{aligned} P([1.35; 1.45] \times \{2.25\}) &\subseteq ([1.404; 1.425], 2.25) \\ DP([1.35; 1.45] \times \{2.25\}) &\subseteq \begin{pmatrix} [-0.4; 0.4] & [-0.08; 0.08] \\ 0 & 0 \end{pmatrix} \end{aligned}$$

One can conclude from the numerical figures for P that the van der Pol system maps the initial set onto itself. The enclosures for DP tell us something about the stability of the apparent limit cycle. In the first row of the result, we see that a deviation in the x direction yields a smaller (at most a factor of 0.4) deviation in the x direction and a deviation in the y direction yields a much smaller (at most a factor of 0.08) deviation in the x direction. Since this Poincaré section is perpendicular to the y axis, there is no deviation in that direction (zeros in the second row of the result).

5 Uniform Limit as Filter

A noteworthy difference to textbook presentations is the way we work with limits. For the results in this article, we needed in particular the notion of uniform limit. In order to define uniform convergence, we use filters. Filters have proved to be useful to describe all kinds of limits and convergence in a coherent way [12]. For details about filters, please consider the paper [12]. In the formalization, the uniform limit *uniform-limit* $X f l F$ is parameterized by a filter F , here we just present the explicit formulations for the *sequentially* and *at* filters.

A sequence of functions $f_n : \alpha \rightarrow \beta$ for $n \in \mathbb{N}$ is said to converge uniformly on $X : \mathcal{P}(\alpha)$ against the *uniform limit* $l : \alpha \rightarrow \beta$, if

Definition 6

$$\begin{aligned} \text{uniform-limit } X f l \text{ sequentially} &:= \\ \forall \varepsilon > 0. \exists N. \forall x \in X. \forall n \geq N. \text{dist}(f_n x) (l x) &< \varepsilon \end{aligned}$$

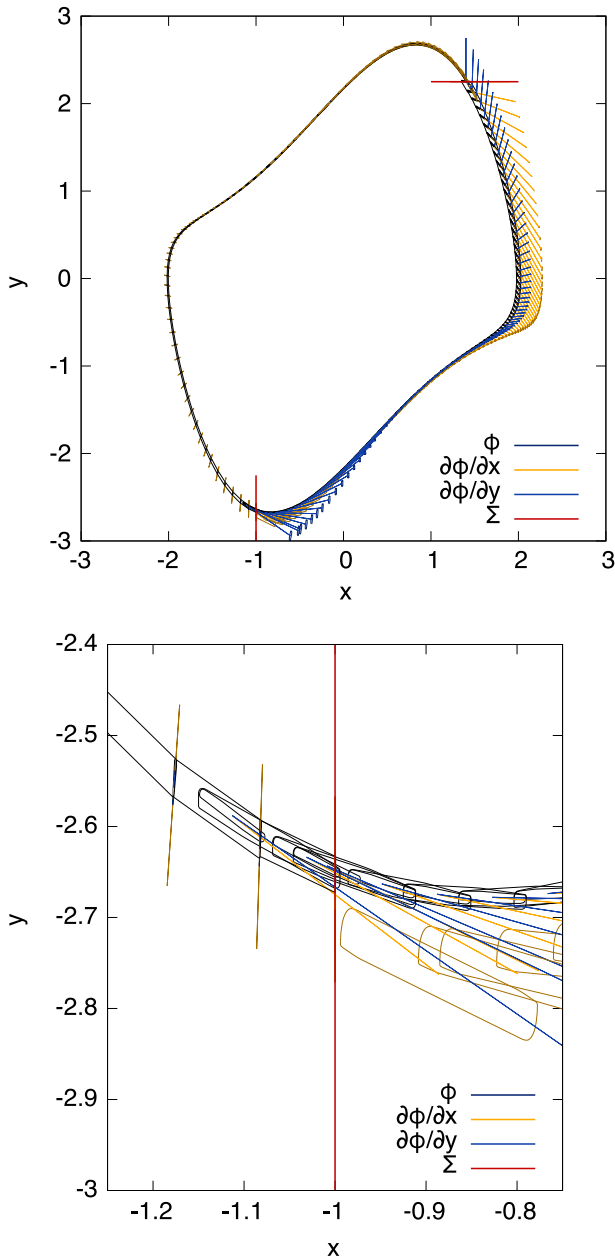


Fig. 3 Computation of Poincaré map $P([1.35; 1.45] \times \{2.25\})$ of the van der Pol system $(\dot{x}, \dot{y}) = (y, (1 - x^2)y - x)$ with Poincaré section $\Sigma = [1; 2] \times \{2.25\}$ and reduction at $\{-1\} \times [1; 2]$ (detailed in lower part). Dark blue and dark yellow lines enclose uncertainties in the numeric values

Note the difference to point-wise convergence, where one would exchange the order of the quantifiers $\exists N$ and $\forall x \in X$.

With the *(at z)* filter, we can also handle uniform convergence of a family of functions $f_y : \alpha \rightarrow \beta$ as y approaches z :

Definition 7

$$\text{uniform-limit } X f l \text{ (at } z) := \forall \varepsilon > 0. \exists \delta > 0. \forall y. |y - z| < \delta \implies (\forall x \in X. \text{dist}(f_y x) (l x) < \varepsilon)$$

The advantage of the filter approach is that many important lemmas can be expressed for arbitrary filters, for example the uniform limit theorem, which states that the uniform limit l of a (via filter F generalized) sequence f_n of continuous functions is continuous.

Theorem 7 (Uniform Limit Theorem)

$$(\forall n \in F. \text{continuous-on } X f_n) \implies \text{uniform-limit } X f l F \implies \text{continuous-on } X l$$

A frequently used criterion to show that an infinite series of functions converges uniformly is the Weierstrass M-test. Assuming majorants M_n for the functions f_n and assuming that the series of majorants converges, it allows one to deduce uniform convergence of the partial sums towards the series.

Lemma 2 (Weierstrass M-Test)

$$\forall n. \forall x \in X. |f_n x| \leq M_n \implies \sum_{n \in \mathbb{N}} M_n < \infty \implies \text{uniform-limit } X \left(n \mapsto x \mapsto \sum_{i \leq n} f_i x \right) \left(x \mapsto \sum_{i \in \mathbb{N}} f_i x \right) \text{ sequentially}$$

6 Bounded Linear Functions

Function spaces, i.e., sets of functions equipped with a certain structure (e.g. topology or norm) are often required for our formalizations. Textbooks readily introduce them with the appropriate definitions, whereas Isabelle/HOL requires more infrastructure. This is due to the fact that the hierarchy of topological spaces is formalized using type classes. In order to profit from this formalization, one needs to introduce types for such spaces. In this section, we introduce a type of (bounded) linear functions, its instantiation as a normed vector space, and how it is used in our formalization.

6.1 Type Classes for Mathematics in Isabelle/HOL

In Isabelle/HOL, many of the mathematical concepts (in particular spaces with a certain structure) are formalized using type classes. The advantage of type class based reasoning is that most of the reasoning is generic: formalizations are carried out in the context of type classes and can then be used for all types inhabiting that type class. For generic formalizations, we use Greek letters α, β, γ and name their type class constraints in prose (i.e., if we write that we “consider a topological space” α , then this result is formalized generically for every type α that fulfills the properties of a topological space).

The spaces we consider are topological spaces with *open* sets, (real) vector spaces with addition $+$: $\alpha \rightarrow \alpha \rightarrow \alpha$ and scalar multiplication $(_) (_) : \mathbb{R} \rightarrow \alpha \rightarrow \alpha$. Normed vector spaces come with a norm $|_| : \alpha \rightarrow \mathbb{R}$. A vector space with multiplication $*$: $\alpha \rightarrow \alpha \rightarrow \alpha$ that is compatible with addition $(a + b) * c = a * c + b * c$ is an algebra and can also be endowed with a norm. Complete normed vector spaces are called Banach spaces.

6.2 A Type of Bounded Linear Functions

An important concept is that of a linear function. For vector spaces α and β , a linear function is a function $f : \alpha \rightarrow \beta$ that is compatible with addition and scalar multiplication.

Definition 8

$$\text{linear } f := \forall x \ y \ c. f(cx + y) = cf(x) + f(y)$$

We need topological properties of linear functions, we therefore now assume normed vector spaces α and β . One usually wants linear functions to be continuous, and if α and β are vector spaces of finite dimension, any linear function $\alpha \rightarrow \beta$ is continuous. In general, this is not the case, and one usually assumes *bounded* linear functions. The norm of the result of a bounded linear function is linearly bounded by the norm of the argument:

Definition 9

$$\text{bounded-linear } f := \text{linear } f \wedge \exists K. \forall x. |f(x)| \leq K|x|$$

We now cast bounded linear functions $\alpha \rightarrow \beta$ as a type $\alpha \rightarrow_{bl} \beta$ in order to make it an instance of Banach space.

Definition 10

$$\text{typedef } \alpha \rightarrow_{bl} \beta := \{f : \alpha \rightarrow \beta \mid \text{bounded-linear } f\}$$

6.3 Instantiations

For defining operations on type $\alpha \rightarrow_{bl} \beta$, the Lifting and Transfer package [13] is an essential tool: operations on the plain function type $\alpha \rightarrow \beta$ are automatically lifted to definitions on the type $\alpha \rightarrow_{bl} \beta$ when supplied with a proof that functions in the result are *bounded-linear* under the assumption that argument functions are *bounded-linear*. We write application of a bounded linear function $f : \alpha \rightarrow_{bl} \beta$ with an element $x : \alpha$ as follows.

Definition 11 (Application of Bounded Linear Functions)

$$(f \cdot x) : \beta$$

We present the definitions of operations involving the type $\alpha \rightarrow_{bl} \beta$ by presenting them in an extensional form using \cdot . Bounded linear functions with pointwise addition and pointwise scalar multiplication form a vector space.

Definition 12 (Vector Space Operations)

For $f, g : \alpha \rightarrow_{bl} \beta$ and $c : \mathbb{R}$,

$$\begin{aligned} (f + g) \cdot x &:= f \cdot x + g \cdot x \\ (cf) \cdot x &:= c(f \cdot x) \end{aligned}$$

The usual choice of a norm for bounded linear functions is the operator norm: the maximum of the image of the bounded linear function on the unit ball. With this norm, $\alpha \rightarrow_{bl} \beta$ forms a normed vector space and we prove that it is Banach if α and β are Banach.

Definition 13 (*Norm in Banach Space*) For $f : \alpha \rightarrow_{bl} \beta$,

$$|f| := \max \{|f \cdot y| \mid |y| \leq 1\}$$

One can also compose bounded linear functions according to $(f \circ g) \cdot x = f \cdot (g \cdot x)$. Bounded linear operators—that is bounded linear functions $\alpha \rightarrow_{bl} \alpha$ from one type α into itself—form a Banach algebra with composition as multiplication and the identity function 1_{bl} as neutral element:

Definition 14 (*Banach Algebra of Bounded Linear Operators*) For $f, g : \alpha \rightarrow_{bl} \alpha$,

$$\begin{aligned} (f * g) \cdot x &:= (f \circ g) \cdot x \\ 1_{bl} \cdot x &:= x \end{aligned}$$

With these instantiations, we can profit from many of the developments that are available for Banach spaces or algebras. In particular, the type of bounded linear functions can be used to describe derivatives in arbitrary vector spaces (Sect. 6.4) and allows one to naturally express e.g., continuity of derivatives, topological properties of inverse linear functions (Sect. 6.5), or the implicit function theorem (Sect. 6.6).

6.4 Total Derivatives

The total derivative (or Fréchet derivative) is a generalization of the ordinary derivative (of functions $\mathbb{R} \rightarrow \mathbb{R}$) for arbitrary normed vector spaces. To illustrate this generalization, recall that the ordinary derivative yields the slope of the function: if $f'(x) = m$, then

$$\lim_{h \rightarrow 0} \frac{f(x+h) - f(x)}{h} = m \tag{2}$$

Moving the m under the limit, one sees that the (linear) function $h \mapsto hm$ is a good approximation for the difference of the function value at nearby points x and $x + h$:

$$\lim_{h \rightarrow 0} \frac{f(x+h) - f(x) - hm}{h} = 0$$

This concept can be generalized by replacing $h \mapsto hm$ with an arbitrary (bounded) linear function A . In the following equation, A is a good linear approximation.

$$\lim_{h \rightarrow 0} \frac{f(x+h) - f(x) - A \cdot h}{|h|} = 0 \tag{3}$$

Note that in the previous equation, we can (just formally) drop many of the restrictions on the type of f . We started with $f : \mathbb{R} \rightarrow \mathbb{R}$ in equation 2, but the last equation still makes sense for $f : \alpha \rightarrow \beta$ for normed vector spaces α, β . We call $A : \alpha \rightarrow_{bl} \beta$ the total derivative Df of f at a point x :

Definition 15 (*Total Derivative*) For $A : \alpha \rightarrow_{bl} \beta$ in equation 3, we write

$$Df|_x = A$$

The total derivative is important for our developments as it is for example the derivative W of the flow in Theorem 4. It is only due to the fact that the resulting type $\alpha \rightarrow_{bl} \alpha$ is a normed vector space, that makes it possible to express continuity of the derivative or to express higher derivatives.

6.5 Inverse Functions

In the Banach space of bounded linear functions, the set of invertible functions is open:

Theorem 8

$$\text{open } \{f :: \alpha \rightarrow_{bl} \beta \mid \exists f^{-1}. f \circ f^{-1} = 1 \wedge f^{-1} \circ f = 1\}$$

The proof of this theorem is based on the fact that (in the Banach Algebra of Bounded Operators), the inverse of the disturbed identity function $1 + w$ with $|w| < 1$ is the convergent series $\sum_i (-1)^i w^i$:

Lemma 3 For $1_{bl}, w :: \alpha \rightarrow_{bl} \alpha$ with $|w| < 1$, $(\sum_i (-1)^i w^i)$ is convergent and the left and right inverse of $1 + w$:

$$\left(\sum_i (-1)^i w^i\right) * (1_{bl} + w) = 1_{bl} \wedge \left(\sum_i (-1)^i w^i\right) * (1_{bl} + w) = 1_{bl}$$

Moreover, one can bound the norm of the inverse,

Lemma 4 $|(1 + w)^{-1} - (1 + w)| \leq \frac{|w|^2}{1 - |w|}$

which is used to prove the set of invertible linear functions open for Theorem 8.

6.6 Implicit Function Theorem

The implicit function theorem is a powerful tool to construct (differentiable) functions satisfying a given (implicit) equation. Given an ‘‘equation’’ $F :: \mathbb{R}^a \times \mathbb{R}^b \rightarrow \mathbb{R}^c$ with a root $F(x, y) = 0$, the theorem allows to ‘‘extend’’ the root (x, y) in an ε -neighborhood $U_\varepsilon(x)$ to a solution function u . This means that the image of the graph of u under F remains constant, namely $F(x, u(x)) = 0$.

For a precise formulation, F needs to be continuously differentiable with invertible derivative:

Theorem 9 (Implicit Function Theorem) *Assume a zero $F(x, y) = 0$ of a continuously differentiable function F . We use the following notation for the derivative of F w.r.t. the 1st argument $f_1 \cdot d := \mathcal{D}F|_{(x,y)}(d, 0)$ and the derivative of F w.r.t. the 2nd argument $f_2 \cdot d := \mathcal{D}F|_{(x,y)}(0, d)$. Assume that f_2 is invertible, i.e., f_2^{-1} exists with $f_2^{-1} \circ f_2 = 1_{bl}$ and $f_2 \circ f_2^{-1} = 1_{bl}$.*

Then there exist u and $\varepsilon > 0$ such that:

- $u(x) = y$
- $F(x, u(x)) = 0$
- $\forall s \in U_\varepsilon(x). F(s, u(s)) = 0$
- continuous-on $U_\varepsilon(x)$ u
- $Du|_x = -f_2^{-1} \circ f_1$
- u is unique: for every v, V where $V \subseteq U_\varepsilon(x)$ is open and connected, v with continuous-on V $v, v(x) = y$, and $(\forall s \in V. F(s, v(s)) = 0)$, it holds that $\forall s \in V. v(s) = u(s)$

Existence of such a function u on a neighborhood $U_\varepsilon(x)$ can be reduced to the *inverse function theorem*, which already exists in Isabelle’s library. We therefore perform this reduction, which yields the expression for the derivative $Du|_x = -f_2^{-1} \circ f_1$. Openness of invertible linear maps (Theorem 8) is required for this construction.

6.7 Further Examples

Here we illustrate in some (for this article only tangentially relevant) examples, that the type of bounded linear functions is useful to conveniently formalize some basic results from analysis: the Leibniz rule for differentiation under the integral sign and conditions for (total) differentiability of multidimensional functions. Furthermore, the exponential function is defined generically for *banach-algebra* and can therefore be used for bounded linear functions as well.

6.7.1 Exponential of Operators

The exponential function for bounded linear functions is a useful concept and important for the analysis of linear ODEs. Here we present that the solution of linear autonomous homogeneous differential equations can be expressed using the exponential function. For a Banach algebra α , the exponential function is defined using the usual power series definition (B^k is a k fold multiplication $B * \dots * B$):

Definition 16 (*Exponential Function*) For a Banach algebra α and $B : \alpha$,

$$e^B := \sum_{k=0}^{\infty} \frac{1}{k!} B^k$$

We prove the following rule for the derivative of the exponential function

Lemma 5 (*Derivative of Exponential*) $\frac{d e^{tA}}{dt} = e^{tA} A$

Proof After unfolding the definition of derivative $\frac{d e^{tA}}{dt} = \lim_{h \rightarrow 0} \frac{e^{(t+h)A} - e^{tA}}{h}$, the crucial step in the proof is to exchange the two limits (one is explicit in $\lim_{h \rightarrow 0}$, and the other one is hidden as the limit of the series Definition 16 of the exponential). Exchange of limits can be done similar to Theorem 7, while uniform convergence is guaranteed according to the Weierstrass M-Test from Lemma 2. □

With this rule for the derivative and an obvious calculation for the initial value, one can show the following

Lemma 6 (*Solution of linear initial value problem*) $\phi_{x_0, t_0}(t) := (e^{(t-t_0)A})(x_0)$ is the unique solution to the ODE $\dot{\phi} t = A(\phi t)$ with initial condition $\phi(t_0) = x_0$.

6.7.2 Total Derivative via Continuous Partial Derivatives

Another example, where interpreting the derivative as bounded linear function $\alpha \rightarrow_{bl} \beta$ is helpful, is when deducing the total derivative of a function f by looking at its partial derivatives f_1 and f_2 (that is, the derivatives w.r.t. one variable while fixing the other). One needs the assumption that the partial derivatives are continuous.

Lemma 7 (*Total Derivative via Continuous Partial Derivatives*) For $f : \alpha \rightarrow \beta \rightarrow \gamma$, $f_1 : \alpha \rightarrow \beta \rightarrow (\alpha \rightarrow_{bl} \gamma)$, $f_2 : \alpha \rightarrow \beta \rightarrow (\beta \rightarrow_{bl} \gamma)$

$$\begin{aligned} \forall x. \forall y. D(x \mapsto f x y)|_x &= f_1 x y \implies \\ \forall x. \forall y. D(y \mapsto f x y)|_y &= f_2 x y \implies \\ \text{continuous}((x, y) \mapsto f_1 x y) &\implies \\ \text{continuous}((x, y) \mapsto f_2 x y) &\implies \\ D((x, y) \mapsto f x y)|_{(x,y)} \cdot (t_1, t_2) &= (f_1 x y) \cdot t_1 + (f_2 x y) \cdot t_2 \end{aligned}$$

6.7.3 Leibniz Rule

Another example is a general formulation of the Leibniz rule. The following rule is a generalization of e.g., the rule formalized by Lelay and Melquiond [20] to general vector spaces. Here $[[a; b]]$ is a hyperrectangle in Euclidean space \mathbb{R}^n . The rule allows one to differentiate under the integral sign: the derivative of the parameterized integral $\int_a^b f(x, t) dt$ with respect to x can be expressed as the integral of the derivative of f . Note that the integral on the right is in the Banach space of bounded linear functions.

Lemma 8 (Leibniz rule) *For Banach spaces α, β and $f : \alpha \rightarrow \mathbb{R}^n \rightarrow \beta, f_1 : \alpha \rightarrow \mathbb{R}^n \rightarrow (\alpha \rightarrow_{bl} \beta)$,*

$$\begin{aligned} \forall t. D(x \mapsto f(x, t))|_x &= f_1(x, t) \implies \\ \forall x. (f(x, \cdot)) \text{ integrable-on } [[a; b]] &\implies \\ \text{continuous-on } (\alpha \times [[a; b]]) ((x, t) \mapsto f(x, t)) &\implies \\ D\left(x \mapsto \int_a^b f(x, t) dt\right)|_x &= \int_a^b f_1(x, t) dt \end{aligned}$$

7 Details About the Flow

We will now go into the technical details of the proofs leading towards continuity and differentiability of the flow (Theorems 3 and 4). We still do not present the proofs: their structure is very similar to the textbook [11] proofs. Nevertheless, we want to present the detailed statements of the propositions, as they give a good impression on the kind of reasoning that was required.

7.1 Criteria for Unique Solution

First of all, we specify the common assumptions to guarantee existence of a unique solution for an initial value problem and therefore a condition for the flow in Definition 1 to be well-defined.

We assume that f is locally Lipschitz continuous in its second argument. That is, for every $(t, x) \in T \times X$ there exist ε -neighborhoods $U_\varepsilon(t)$ and $U_\varepsilon(x)$ around t and x , in which f is Lipschitz continuous w.r.t. the second argument (uniformly w.r.t. the first: the L is valid for all t').

Definition 17

$$\begin{aligned} \text{local-lipschitz } T \ X \ f &:= \\ \forall t \in T. \forall x \in X. & \\ \exists \varepsilon > 0. \exists L. & \\ \forall t' \in U_\varepsilon(t). \forall x_1, x_2 \in U_\varepsilon(x). & |f(t', x_1) - f(t', x_2)| \leq L|x_1 - x_2| \end{aligned}$$

Now the only assumptions that we need to prove continuity of the flow are open sets for time and phase space and a locally Lipschitz continuous right-hand side f that is continuous in t .

Definition 18 (Conditions for unique solution)

1. T is an open set

2. X is an open set
3. f is locally Lipschitz continuous on X : *local-lipschitz* $T X f$
4. for every $x \in X, t \mapsto f(t, x)$ is continuous on T .

The detailed proofs that these assumptions guarantee the existence of a unique solution for initial value problems have been presented in Theorem 3 of earlier work [17].

7.2 The Frontier of the State Space

It is important to study the behavior of the flow at the frontier of the state space (e.g., as time or the solution tend to infinity). From this behavior, one can deduce conditions under which solutions can be continued. This yields techniques to gain more precise information on the existence interval *ex-ivl*.

If the solution only exists for finite time, it has to explode (i.e., leave every compact set):

Lemma 9 (Explosion for Finite Existence Interval)

$$\text{ex-ivl}(x_0) =]\alpha, \beta[\implies \beta < \infty \implies \text{compact } K \implies \\ \exists t \geq 0. t \in \text{ex-ivl}(x_0) \wedge \phi(x_0, t) \notin K$$

This lemma can be used to prove a condition on the right-hand side f of the ODE, to certify that the solution exists for the whole time. Here the assumption guarantees that the solution stays in a compact set.

Lemma 10 (Global Existence of Solution)

$$(\forall s \in T. \forall u \in T. \exists L. \exists M. \forall t \in [s; u]. \forall x \in X. |f t x| \leq M + L|x|) \\ \implies \text{ex-ivl}(x_0) = T$$

7.3 Continuity of the Flow

The following lemmas are all related to continuity of the flow. With the help of Grönwall’s Lemma 1, one can show that when two solutions (starting from different initial conditions x_0 and y_0) both exist for a time t and are restricted to some set Y on which the right-hand side f satisfies a (global) Lipschitz condition K , then the distance between the solutions grows at most exponentially with increasing time:

Lemma 11 (Exponential Initial Condition for Two Solutions)

$$t \in \text{ex-ivl}(x_0) \implies t \in \text{ex-ivl}(y_0) \implies \\ x_0 \in Y \implies y_0 \in Y \implies Y \subseteq X \implies \\ \forall s \in [0; t]. \phi(x_0, s) \in Y \implies \\ \forall s \in [0; t]. \phi(y_0, s) \in Y \implies \\ \forall s \in [0; t]. \text{lipschitz } Y (f s) K \implies \\ |\phi(x_0, t) - \phi(y_0, t)| \leq |x_0 - y_0|e^{K|t|}$$

Note that it can be hard to establish the assumptions of this lemma, in particular the assumption that both solutions from x_0 and y_0 exist for the same time t . Consider Fig. 1: not all solutions (e.g., from y_0) do necessarily exist for the same time s (e.g., the solution from z_0). One can choose, however, a neighborhood of y_0 (e.g., including x_0), such that all solutions starting from within this neighborhood exist for at least the same time, and with

the help of the previous lemma, one can show that the distance of these solutions increases at most exponentially:

Lemma 12 (Exponential Initial Condition of Close Solutions)

$$\begin{aligned}
 a \in \text{ex-ivl}(x_0) &\implies b \in \text{ex-ivl}(x_0) \implies a \leq b \implies \\
 \exists \delta > 0. \exists K > 0. U_\delta(x_0) &\subseteq X \wedge \\
 (\forall y \in U_\delta(x_0). \forall t \in [a; b]. & \\
 t \in \text{ex-ivl}(y) \wedge |\phi(x_0, t) - \phi(y, t)| &\leq |x_0 - y|e^{K|t|})
 \end{aligned}$$

Using this lemma is the key to showing continuity of the flow (Theorem 3).

A different kind of continuity is not with respect to the initial condition, but with respect to the right-hand side of the ODE.

Lemma 13 (Continuity with respect to ODE) *Assume two right-hand sides f, g defined on X and uniformly close $|f(x) - g(x)| < \varepsilon$. Furthermore, assume a global Lipschitz constant K for f on X . Then the deviation of the flows ϕ_f and ϕ_g can be bounded:*

$$|\phi_f(x_0, t) - \phi_g(x_0, t)| \leq \frac{\varepsilon}{K} e^{K|t|}$$

7.4 Differentiability of the Flow

The proof for the differentiability of the flow incorporates many of the tools that we have presented up to now, we will therefore go a bit more into the details of this proof.

7.4.1 Assumptions

The assumptions in Definition 18 are not strong enough to prove differentiability of the flow. However, a continuously differentiable right-hand side $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ suffices. To be more precise:

Definition 19 (Criterion for Continuous Differentiability of the Flow)

$$\exists f' : \mathbb{R}^n \rightarrow (\mathbb{R}^n \rightarrow_{bl} \mathbb{R}^n). (\forall x \in X. \text{Df}|_x = f'(x) \wedge \text{continuous-on } X \text{ } f')$$

From now on, we denote the derivative along the flow from x_0 with $A_{x_0} : \mathbb{R} \rightarrow \mathbb{R}^n$:

Definition 20 (Derivative along the Flow) $A_{x_0}(t) := \text{Df}|_{\phi(x_0, t)}$

The derivative of the flow is the solution to the so-called variational equation, a non-autonomous linear ODE. The initial condition ξ is supposed to be a perturbation of the initial value (like $\frac{\partial \phi}{\partial x}$ and $\frac{\partial \phi}{\partial y}$ in Fig. 2) and in what follows we will prove that the solution to this ODE is a good (linear) approximation of the propagation of this perturbation.

$$\begin{cases} \dot{u}(t) = A_{x_0}(t) \cdot u(t) \\ u(0) = \xi \end{cases}, \tag{4}$$

We will write $u_{x_0}(\xi, t)$ for the flow of this ODE and omit the parameter x_0 and/or the initial value ξ if they can be inferred from the context.

As a prerequisite for the next proof, we begin by proving that $u_{x_0}(\xi, t)$ is linear in ξ , a property that holds because u is the solution of a linear ODE (this is often also called the ‘‘superposition principle’’).

Lemma 14 (Linearity of $u_{x_0}(\xi, t)$ in ξ)

$$\alpha u_{x_0,a}(t) + \beta u_{x_0,b}(t) = u_{x_0,(\alpha a + \beta b)}(t).$$

Because $\xi \mapsto u_{x_0}(\xi, t) : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is linear on Euclidean space, it is also bounded linear, so we will identify this function with the corresponding element of type $\mathbb{R}^n \rightarrow_{bl} \mathbb{R}^n$. The main efforts go into proving the following lemma, showing that the aforementioned function is the derivative with respect to space of the flow for fixed time $\phi_t = (x_0 \mapsto \phi(x_0, t))$.

Lemma 15 (Space Derivative of the Flow) For $t \in \text{ex-ivl}(x_0)$,

$$(D\phi_t|_{x_0}) \cdot \xi = u_{x_0}(\xi, t)$$

Proof The proof starts out with the integral identities of the flow, the perturbed flow, and the linearized propagation of the perturbation:

$$\begin{aligned} \phi(x_0, t) &= x_0 + \int_0^t f(\phi(x_0, s)) \, ds \\ \phi(x_0 + \xi, t) &= x_0 + \xi + \int_0^t f(\phi(x_0 + \xi, s)) \, ds \\ u_{x_0}(\xi, t) &= \xi + \int_0^t A_{x_0}(s) \cdot u_{x_0}(\xi, s) \, ds \\ &= \xi + \int_0^t f'(\phi(x_0, s)) \cdot u_{x_0}(\xi, s) \, ds \end{aligned}$$

Then, for any fixed ε , after a sequence of estimations (3 pages in the textbook proof, about 500 lines of proof text) involving e.g., uniform convergence (Sect. 5) of the first-order remainder term of the Taylor expansion of f , continuity of the flow (Theorem 3), and linearity of u (Lemma 14) one can prove the following inequality.

$$\frac{\|\phi(x_0 + \xi, t) - \phi(x_0, t) - u_{x_0}(\xi, t)\|}{\|\xi\|} \leq \varepsilon$$

This shows that $u_{x_0}(\xi, t)$ is indeed a good approximation for the propagation of the initial perturbation ξ and exactly the definition for the space derivative of the flow. □

Note that $u_{x_0}(\xi, t)$ yields the space derivative in direction of the vector ξ . The total space derivative of the flow is then the linear function $\xi \mapsto u_{x_0,\xi}(t)$. But this derivative can also be described as the solution of the following “matrix-valued” variational equation:

$$\begin{cases} \dot{W}_{x_0}(t) = A_{x_0}(t) \circ W_{x_0}(t) \\ W_{x_0}(0) = \text{Id} \end{cases} \tag{5}$$

This initial value problem is defined for linear operators of type $\mathbb{R}^n \rightarrow_{bl} \mathbb{R}^n$. Thanks to Lemma 10, one can show that it is defined on the same existence interval as the flow ϕ . The solution W_{x_0} is related to solutions of the variational IVP as follows:

$$u_{x_0}(\xi, t) = W_{x_0}(t) \cdot \xi$$

The derivative of the flow ϕ at (x_0, t) with respect to t is given directly by the ODE, namely $f(\phi(x_0, t))$. Therefore and according to Lemma 7 the total derivative of the flow is characterized as follows:

Theorem 10 (Derivative of the Flow)

$$D\phi|_{(x_0,t)} \cdot (\xi, \tau) = W_{x_0}(t) \cdot \xi + \tau f(\phi(x_0, t))$$

7.5 Continuity of Derivative

Regarding the continuity of the derivative $D\phi|_{(x_0,t)} \cdot (\xi, \tau)$ with respect to (x_0, t) : $\tau f(\phi(x_0, t))$ is continuous because of Definition 18 and Theorem 3.

$W_{x_0}(t)$ is continuous with respect to t , so what remains to be shown is continuity of the space derivative regarding x_0 . The proof of this statement relies on Theorem 13, because for different values of x_0 , W_{x_0} is the solution to ODEs with slightly different right-hand sides. A technical difficulty here is to establish the assumption of *global* Lipschitz continuity for Theorem 13.

8 Details About the Poincaré Map

Here we sketch how to prove continuity and differentiability of the Poincaré map. We assume a Poincaré section to be a subset (via S) of an implicit surface:

$$\Sigma = \{x \in S \mid s(x) = 0\}$$

The idea is to apply the implicit function Theorem 9 to find a differentiable function u that solves $s(\phi(x, u(x))) = 0$ in a neighborhood $U_\varepsilon(x)$. The implicit function u is unique (w.r.t. continuous functions). Since $s(\phi(x, \tau(x))) = 0$, all one has to do is prove that τ is continuous, in order to get differentiability for τ , because τ is equal to u and therefore $D\tau|_x = Du|_x$.

This is the most prominent difference to the construction in textbooks (e.g., [24, chapter 3], [25, chapter 5.8], [11, chapter 10.3]). Those textbooks simply *declare* the return time τ to be the solution u from the implicit function theorem. We use the uniqueness condition of the implicit function Theorem 9 and results about continuity (Theorems 11 and 12) to show the equality of (our global notion of return time) τ and the solution of the implicit function theorem u .

But first, one needs more assumptions on the Poincaré section Σ in order to carry out the construction of the implicit function u : we assume S closed and s continuously differentiable, moreover the flow needs to be transversal at the return map ($Ds|_{P(x)} \cdot f(P(x)) \neq 0$). Moreover, the return $P(x)$ needs to be in the relative (with respect to the implicit surface) interior of the Poincaré section ($\exists \delta. U_\delta(P(x)) \cap \{x \mid s(x) = 0\} \subseteq S$). We summarize these assumptions in the following definition.

Definition 21 (*Assumptions for Poincaré Section*)

$$Ds|_{\phi(x,\tau(x))} \cdot f(P(x)) \neq 0 \wedge (\exists \delta. U_\delta(P(x)) \cap \{x \mid s(x) = 0\} \subseteq S)$$

There are two cases for which we prove $\tau(x)$ continuous: First, if $x \notin \Sigma$, then τ is continuous in any sufficiently small neighborhood around x . The limit can therefore be approached from an arbitrary set X , i.e., *at x within X* . Second, if $x \in \Sigma$, then τ is continuous only the side of the surface Σ to which the vector field points to, i.e., *at x within $\{x \mid s(x) \leq 0\}$* . That is because if y is taken (arbitrarily) close to x , but on the other side of Σ , then $\tau(y)$ is (arbitrarily) close to zero, but $\tau(x) > 0$. More formally, the two cases are given by Theorems 11 and 12:

Theorem 11 (Continuity of τ outside Σ) *continuous τ (at x within X), if the assumptions from Definition 21 and the following conditions hold:*

- *returns-to* (Σ, x)
- $x \notin \Sigma$

Theorem 12 (Continuity of τ on Σ) *continuous τ (at x within $\{x \mid s(x) \leq 0\}$), if the assumptions from Definition 21 and the following conditions hold:*

- *returns-to* (Σ, x)
- $x \in \Sigma$
- $Ds|_x \cdot f(x) < 0$

When τ is continuous (e.g., under the same assumptions as those of Theorem 11), τ is equal to the solution of the implicit function theorem. This yields an explicit expression for the derivative of $D\tau$.

Theorem 13 (Derivative of τ) *Under the assumptions of Theorem 11:*

$$D\tau|_x \cdot h = -\frac{Ds|_{P(x)} \cdot (D\phi_{\tau(x)}|_x \cdot h)}{Ds|_{P(x)} \cdot (f(P(x)))} f(P(x))$$

From continuity and differentiability of the flow (Theorems 3 and 4) and τ (Theorems 11, 12, and 13), it follows that the Poincaré map is continuous and differentiable.

9 Related Work

We first concentrate on related work that has directly influenced the emergence of the mathematics library in Isabelle on which we build our formalization before discussing related work on ODEs in other proof assistants. An extensive survey on real analysis in proof assistants is given by Boldo et al. [4].

Isabelle’s mathematics library has its origins in Fleuriot and Paulson’s [5] theory of real analysis which was mostly specific to the type \mathbb{R} . This material has since been generalized to type classes. Much of the formalization of analysis has been ported to Isabelle/HOL from Harrison’s multivariate analysis library for HOL Light [9, 10]. Most of the theory about derivatives and (Henstock–Kurzweil) integration originates from Harrison’s library, e.g., also the inverse function we refer to in the proof of the implicit function Theorem 9.

Spitters and Makarov [22] implement Picard iteration to compute solutions of ODEs in the interactive theorem prover Coq. Their algorithm yields solutions only for relatively short time intervals, but is a direct result of a constructive proof for existence/uniqueness.

Maggesi [21] used his theory of metric spaces (as a predicate instead of type classes) for a formalization of a local version of the Picard–Lindelöf theorem in HOL Light.

Boldo et al. [2] approximate the solution of one particular partial differential equation with a C-program and verify its correctness in Coq. Another important result, fundamental for numerical approximations of partial differential equations is the formalization of the Lax–Milgram theorem in Coq by Boldo et al. [1]. Their formalization is similar to ours in the sense that it is also based on functional analysis and formalization of bounded linear functions. They build on the Coquelicot library for real analysis [3].

More functional analysis is given by Gouëzel’s formalization of L^p -spaces [7]. Gouëzel [6] also formalized ergodic theory, i.e., dynamical systems with an invariant measure in Isabelle/HOL, his formalization includes Kac’s Formula and Birkhoff and Kingman theorems.

10 Conclusion

To conclude, our formalization of flow and the variational equation contains essentially all lemmas and proofs of at least 22 pages (Chapter 17) of the textbook by Hirsch et al. [11]. This corresponds to Sections 2.1–2.5 on “Nonlinear Systems: Local Theory” (about 30 pages) of Perko’s textbook [24]. The formalization moreover comprises a notion of Poincaré map, which is more flexible than what is usually presented in textbooks.

All of this required general-purpose background to be formalized, in particular uniform limits and the Banach space of (bounded) linear functions. The separate type for bounded linear functions is a minor complication that is necessary to profit from the type class based library for analysis in Isabelle/HOL. We showed the concrete usability of our results by verifying the connection of the abstract formalization with concrete rigorous numerical algorithms.

Our formalization encompasses fundamental properties (continuity and differentiability) of flow and Poincaré map. This is enough material for a specification of the numerical part of Tucker’s proof, where rigorous enclosures of Poincaré maps and their derivatives are computed. We lack a complete formalization of dependence on parameters (we only formalized uniform continuity according to Lemma 13) and higher derivatives. We have not formalized any deeper theory about dynamical systems. The next interesting steps would be e.g., the Poincaré–Bendixson theorem, the stable manifold theorem, or the Hartmann–Grobmann theorem.

Acknowledgements We would like to thank Professor Dr. Martin Brokate for supervising part of this work as an “interdisciplinary project”. Johannes Hölzl’s suggestions related to filters were very helpful. We would also like to thank the anonymous reviewers for all their suggestions and comments. Part of this work was supported by DFG RTG 1480 and DFG NI 491/16-1.

References

1. Boldo, S., Clément, F., Faissole, F., Martin, V., Mayero, M.: A Coq formal proof of the Lax–Milgram theorem. In: Proceedings of the 6th ACM SIGPLAN Conference on Certified Programs and Proofs, ACM, pp. 79–89 (2017)
2. Boldo, S., Clément, F., Filiâtre, J.C., Mayero, M., Melquiond, G., Weis, P.: Wave equation numerical resolution: a comprehensive mechanized proof of a C program. *J. Autom. Reason.* **50**(4), 423–456 (2013). <https://doi.org/10.1007/s10817-012-9255-4>
3. Boldo, S., Lelay, C., Melquiond, G.: Coquelicot: a user-friendly library of real analysis for Coq. *Math. Comput. Sci.* **9**(1), 41–62 (2015). <https://doi.org/10.1007/s11786-014-0181-1>
4. Boldo, S., Lelay, C., Melquiond, G.: Formalization of real analysis: a survey of proof assistants and libraries. *Math. Struct. Comput. Sci.* **26**(7), 1196–1233 (2016). <https://doi.org/10.1017/S0960129514000437>
5. Fleuriot, J.D., Paulson, L.C.: Mechanizing nonstandard real analysis. *LMS J. Comput. Math.* **3**, 140–190 (2000). <https://doi.org/10.1112/S1461157000000267>
6. Gouezel, S.: Ergodic theory. Archive of formal proofs. http://isa-afp.org/entries/Ergodic_Theory.shtml (2015)
7. Gouezel, S.: Lp spaces. Archive of formal proofs. <http://isa-afp.org/entries/Lp.shtml> (2016)
8. Haftmann, F.: Code generation from specifications in higher-order logic. Dissertation, Technische Universität München, München (2009)
9. Harrison, J.: A HOL theory of Euclidean space. In: J. Hurd, T. Melham (eds.) 18th International Conference on Theorem Proving in Higher Order Logics, TPHOLs, LNCS, vol. 3603, pp. 114–129 (2005)
10. Harrison, J.: The HOL light theory of Euclidean space. *J. Autom. Reason.* **50**(2), 173–190 (2013). <https://doi.org/10.1007/s10817-012-9250-9>

11. Hirsch, M.W., Smale, S., Devaney, R.L.: *Differential Equations, Dynamical Systems, and an Introduction to Chaos*. Elsevier, New York (2013)
12. Hölzl, J., Immler, F., Huffman, B.: Type classes and filters for mathematical analysis in Isabelle/HOL. In: *International Conference on Interactive Theorem Proving*, Springer, pp. 279–294 (2013)
13. Huffman, B., Kunčar, O.: Lifting and transfer: a modular design for quotients in Isabelle/HOL. In: *International Conference on Certified Programs and Proofs*, Springer, pp. 131–146 (2013)
14. Immler, F.: Formally verified computation of enclosures of solutions of ordinary differential equations. In: Badger, J., Rozier, K. (eds.) *NASA Formal Methods*, LNCS, vol. 8430, pp. 113–127. Springer, Berlin (2014)
15. Immler, F.: A verified algorithm for geometric zonotope/hyperplane intersection. In: *Proceedings of the 2015 Conference on Certified Programs and Proofs, CPP '15*, ACM, New York, NY, USA, pp. 129–136 (2015)
16. Immler, F.: Verified reachability analysis of continuous systems. In: Baier, C., Tinelli, C. (eds.) *Tools and Algorithms for the Construction and Analysis of Systems*, vol. 9035, pp. 37–51. Springer, Berlin (2015). https://doi.org/10.1007/978-3-662-46681-0_3
17. Immler, F., Hölzl, J.: Numerical analysis of ordinary differential equations in Isabelle/HOL. In: Beringer, L., Felty, A. (eds.) *Interactive Theorem Proving*, LNCS, vol. 7406, pp. 377–392. Springer, Berlin (2012)
18. Immler, F., Hölzl, J.: Ordinary differential equations. *Archive of Formal Proofs*. Formal Proof Development (2017). http://isa-afp.org/entries/Ordinary_Differential_Equations.html
19. Immler, F., Traut, C.: The flow of ODEs. In: *International Conference on Interactive Theorem Proving*, Springer, pp. 184–199 (2016)
20. Lelay, C., Melquiond, G.: Différentiabilité et intégrabilité en Coq. application à la formule de d'Alembert. In: *JFLA—Journées Francophone des Langages Applicatifs*, 2012. Carnac, France (2012). <https://hal.inria.fr/hal-00642206>
21. Maggesi, M.: A formalization of metric spaces in HOL light. *J. Autom. Reason.* (2017). 10.1007/s10817-017-9412-x
22. Makarov, E., Spitters, B.: The Picard algorithm for ordinary differential equations in Coq. In: Blazy, S., Paulin-Mohring, C., Pichardie, D. (eds.) *Interactive Theorem Proving*, LNCS, vol. 7998, pp. 463–468. Springer, Berlin (2013)
23. Nipkow, T., Paulson, L.C., Wenzel, M.: *Isabelle/HOL: A Proof Assistant for Higher-Order Logic*. Springer, Berlin (2002)
24. Perko, L.: *Differential Equations and Dynamical Systems*. Springer, Berlin (2001). <https://doi.org/10.1007/978-1-4613-0003-8>
25. Robinson, C.: *Dynamical Systems, Stability, Symbolic Dynamics, and Chaos*. CRC Press, Boca Raton (1999). <https://doi.org/10.1007/978-1-4613-0003-8>
26. Tucker, W.: A rigorous ODE solver and Smale's 14th problem. *Found. Comput. Math.* **2**(1), 53–117 (2002)