

AD-like techniques

in

global optimization

A. Neumaier
University of Vienna
Austria

Reliable global optimization

requires, not only access to function values, gradients and Hessians of a function at given points, but also a variety of properties conveying

global information

that can be constructed in an AD-like fashion from a computational graph:

- ranges of functions and constraints
- linear and nonlinear relaxations
- interval slopes
- etc.

Part of what we discovered in this context may also be useful in other contexts, such as

- solutions of nonlinear systems
- local optimization
- algorithm-independent modeling

In this talk, I'll emphasize the latter aspects.

Overview.

1. Slopes

- they generalize derivatives and AD
- they are more robust than derivatives, when applicable

2. FMathL

- A modeling and documentation language for mathematics, under development at our institute in Vienna

3. AD-like propagation methods

- Taylor series and generalizations
- Linear relaxations
- Bounds on variables
- and much else ...

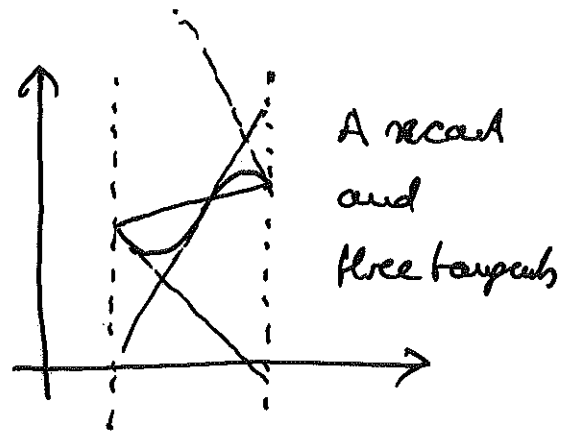
4. Global optimization on graphs

5. Inclusion- and Exclusion regions for zeros of systems of equations

(if time permits)

Slopes I

Slopes are a generalization of derivatives with much better global approximation properties.



There are problems whose definition already requires derivatives, such as bifurcation problems, finding stationary points, etc. Here derivatives are indispensable.

But often, slopes can be used in place of derivatives, and are much more robust in certain contexts.

Slopes can be computed cheaply by AD, and at the same cost as derivatives. + probably easy to adapt

Essentially all properties of derivatives have an extension to slopes. Proofs of these properties are even simpler (since no limit is needed).

Slopes show their real power in global problems:

- nonexistence of a solution
- existence proofs
- inclusion and exclusion regions

Slopes II

Slopes generalize well-known formulas such as

$$x^2 - z^2 = (x+z)(x-z)$$

$$\sqrt{x} - \sqrt{z} = \frac{x-z}{\sqrt{x} + \sqrt{z}}$$

to general multivariate, vector-valued functions $F: \mathbb{R}^n \rightarrow \mathbb{R}^m$

A slope at $z \in \mathbb{R}^n$ is a matrix $F[z, x] \in \mathbb{R}^{m \times n}$ such that

$$F(x) = F(z) + F[z, x](x-z)$$

If F is continuously differentiable, one usually imposes the requirement of continuity. Then

$$F'(x) = F[x, x]$$

is the limiting slope at two equal arguments.

$$\begin{array}{lll} f(x) = x^2 & f[z, x] = x+z & f[x, x] = 2x = f'(x) \\ f(x) = \sqrt{x} & f[z, x] = \frac{1}{\sqrt{x} + \sqrt{z}} & f[x, x] = \frac{1}{2\sqrt{x}} = f'(x) \end{array}$$

The mean value theorem implies, that (for c.d. F)

$$F[z_1, x] := \int_0^1 dt F'(z_1 + t(x-z_1))$$

is a slope; but this is useless in practice.

For $n > 1$, slopes are not unique:

$$x_1 x_2 - z_1 z_2 = (x_2, z_1) \begin{pmatrix} x_1 - z_1 \\ x_2 - z_2 \end{pmatrix} = (z_2, x_1) \begin{pmatrix} x_1 - z_1 \\ x_2 - z_2 \end{pmatrix}$$

So $f(x) = x_1 x_2 \Rightarrow F[x, z] = (x_2, z_1)$ or (z_2, x_1)
are essentially different slopes

Slopes III

Slopes by finite differences:

$$F[z, x] = \frac{F(x) - F(z)}{z - x} \quad \text{for } u=1$$

requires 2 function values

(for $u > 1$ similar formulas with $u+1$ function values)

But:

Numerically unstable when $z \approx x$

\Rightarrow not recommended

Slopes by symbolic manipulators:

by example: $F(x) = x + \sqrt{2x+1}$

$$\begin{aligned} F(x) - F(z) &= x - z + \sqrt{2x+1} - \sqrt{2z+1} \\ &= x - z + \frac{(2x+1) - (2z+1)}{\sqrt{2x+1} + \sqrt{2z+1}} \end{aligned}$$

$$\Rightarrow F[z, x] = 1 + \frac{2}{\sqrt{2x+1} + \sqrt{2z+1}}$$

numerically stable

But:

expressions are unwieldy when F is complicated
(like in symbolic differentiation, which results for $z=x$)

\Rightarrow not recommended

Slopes by AD-like approach:

recommended

stable

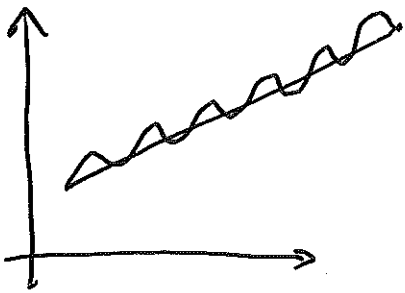
cheap (≤ 3 # Ops for f , for any u)

But needs some preparation... (we first discuss uses, however!)

Slopes IV

Why slopes?

- accurate at 2 points rather than only one
⇒ better global properties
- provides identities rather than truncated series
⇒ full control over errors
- has a smoothing effect
⇒ insensitive to small but high-frequency oscillations or discontinuities
(as frequent in numerical adaptive codes)
whereas derivatives become completely meaningless!



Example $f(x) = e^{i\omega x}$
 $f[x; x] = \frac{e^{i\omega x} - e^{i\omega z}}{x - z} = e^{i\omega z} \frac{e^{i\omega(x-z)} - 1}{x - z}$
 $= e^{i\omega z} i\omega \exp_1(i\omega(x-z))$

where $\exp_1(t) = \frac{e^t - 1}{t}$ is analytic
(and easy to evaluate in terms of $\text{Re}(t)$ approach
for small t , directly for large t)

$$\exp_1(0) = 1 \Rightarrow f[x; x] = e^{i\omega x} i\omega = \mathcal{O}(\omega)$$

$$f[x; x] = \mathcal{O}(\omega) \text{ for } x - z = \mathcal{O}(\omega^{-1}) \text{ only}$$
$$f[x; x] = \mathcal{O}(h^{-1}) \text{ for } x - z = \mathcal{O}(h), \omega^{-1} \ll h$$

Thus large frequencies ($\omega \gg 1$) affect the derivative
but not the slope (except at very short distances where it
is appropriate)

Slopes V

The secant method for $F(x)=0$ in \mathbb{R}^n

Define $N_z(x) := z - F[z, x]^{-1} F(z)$

If $F(\hat{x})=0$ then $F(z) = F(\hat{x}) - F[z, \hat{x}](\hat{x} - z)$
 $= F[z, \hat{x}](z - \hat{x})$

Therefore $F(\hat{x})=0 \Rightarrow N_z(\hat{x}) = \hat{x}$

If $z \approx \hat{x}$ then $N_z(x)$ is a contraction operator near \hat{x}

\Rightarrow Secant method $x_{n+1} = x_n - F[x_n, x_{n-1}]^{-1} F[x_n]$

converges superlinearly (near \hat{x})

(with AD!) cost = cost of Newton's method
speed = slightly less (order 1.6 not 2)

but robustness increased

In practice, needs line search along $p_n = -F[x_n, x_{n-1}]^{-1} F[x_n]$

sparsity = same as Newton's method

Slopes VI

Quasi-Newton updates for optimization using slopes

For a quadratic function $f(x) = \gamma + c^T x + \frac{1}{2} x^T G x$
(G symmetric)

The gradient is $f'(x) = c + Gx$

Write the slope as $f[x_k, x_{k-1}] = (c + \frac{1}{2} G(x_k + x_{k-1}))^T$

So we have the quadratic model

$$f(x) = f(x_k) + f[x_k, x_{k-1}](x - x_k) + \frac{1}{2} (x - x_{k-1})^T G (x - x_k) + O\left(\frac{\|x - x_k\| \|x - x_{k-1}\|}{\|x - x_{k-1}\|}\right)$$

where the error term is just a guess for nonquadratic function (after enough updates)

From two successive slopes one can calculate

the quasi-Newton condition

$$G(x_k - x_{k-2}) = 2(f[x_k, x_{k-1}] - f[x_{k-1}, x_{k-2}])^T$$

so that the Hessian of a quadratic function can be fully accumulated from $n+2$ linearly independent search steps, using the standard SR1 update (and a trust region approach)

The quasi-Newton step from x_k becomes the minimum of the quadratic model, which is at

$$x_{k+1} = \frac{x_k + x_{k-1}}{2} - G^{-1} f[x_k, x_{k-1}]^T$$

if the trust region is not active

- no implementation yet
- finite termination on quadratics \Rightarrow probably superlinear convergence
- should be insensitive to oscillations

Slopes VII

composition formulas for slopes

$$F = \alpha G \pm \beta H \Rightarrow F[z, x] = \alpha G[z, x] \pm \beta H[z, x]$$

$$F = GH \Rightarrow F[z, x] = G[z, x]H(z) + G(x)H[z, x]$$

if G or H is scalar

$$F = G/H \Rightarrow F[z, x] = (G[z, x] - F(z)H[z, x]) / H(x)$$

if H is scalar

$$F = G \circ H \Rightarrow F[z, x] = G[H(z), H(x)] \quad H[z, x] \quad \text{chain rule}$$

(Krawczyk & Neumaier 1985)

For multiplication and division, there is a second formula with (x) and (z) interchanged, which (for $n > 1$) typically gives different results. The same holds for the formula for implicit functions

F defined by $G(F(x), H(x)) = 0$ for $x \approx z$

$$\Rightarrow F[z, x] = -G[F(z), F(x)](H(x))^{-1} G(F(z)) [H(z), H(x)] H[z, x]$$

Here $G[u, v](w)$ is the slope of $G(u, v)$ w.r. to u
 $G(u)[v, w]$ " " " " $G(u, v)$ w.r. to v

This formula seems to be new, and implies application to slopes of solutions of linear systems, inverses, and eigenvalue problems, which can all be stated as implicit functions

As observed by Blich (1992), $c^T F[z, x]$ can be propagated cheaply in a backward fashion, using the standard trick of AD for adjoint computations

\Rightarrow slopes of scalar functions at the cost of 3# ops for slopes of arbitrary functions at same cost as derivatives

Slopes VIII

Slopes for initial-value problems

The parameterized IVP

$$\begin{aligned}\dot{u} &= F(u, x) && \text{(generally vector-valued)} \\ u(0) &= u_0(x)\end{aligned}$$

define $u(t)$ as a function of $t \in [0, T]$.

The slope $w = u[\xi, x]$ then satisfies the

variational equation

$$\begin{aligned}\dot{w} &= F[u(\xi), u(x)](\xi) w + F(u(x))[\xi, x] && \text{(and a second alternative formula)} \\ w(0) &= u_0[\xi, x]\end{aligned}$$

which can be used in forward mode to calculate the slope if there are few parameters

To get corresponding adjoint equations, we write

$$\langle C, u \rangle_t = \int_t^T dt c(t)^T u(t) + b^T u(T).$$

To compute the slope of the linear functional $\langle C, u \rangle_0$, we construct functions $v(t), v_0(t)$ such that

$$\langle C, u(x) - u(\xi) \rangle_t = v(t)^T (u(x) - u(\xi)) \Big|_t + v_0(t)^T (x - \xi)$$

Comparing the derivative w.r. to t of both sides we see that we must satisfy

$$c^T(u(x) - u(\xi)) = \dot{v}^T (u(x) - u(\xi)) + v^T (\dot{u}(x) - \dot{u}(\xi)) + \dot{v}_0^T (x - \xi)$$

This is satisfied if the slope satisfies the equation

$$c^T w[\xi, x] = \dot{v}^T w[\xi, x] + v^T \dot{w}[\xi, x] + \dot{v}_0^T$$

Stokes VIII, ctd.

Using the variational equation, we can match the two sides if we solve the coupled adjoint equations (backward in time)

$$\dot{v} = c - F(u(t), u(x)) (z)^T v$$

$$v(T) = b$$

and

$$\dot{v}_0 = -F(u(x)) [z, x]^T v$$

$$v_0(T) = 0$$

(The initial conditions at $t=T$ are obvious.)

Slopes IX

Verifying existence or nonexistence of zeros

$F(x) = 0$, $x \in X$ (a nonempty, closed, convex set, in practice, a box: $X = [x, \bar{x}] \subseteq \mathbb{R}^n$)

We know that any zero \bar{x} is a fixed point of

$$N(x) = z - F[z, x]^{-1} F(z)$$

Using appropriate fixed-point theorems, we can conclude existence of a solution provided that a set theoretic closure condition is satisfied; typically the fixed point operator must map the region of interest into itself, and some other technical conditions must hold.

This requires bounds on ranges over sets, which are generally computed using interval arithmetic

In the simplest case, one just substitutes intervals for the variables and computes an enclosure.

(If the result is pessimistic, a more sophisticated analysis is called for; there are a number of pitfalls that must be recognized and overcome.)

Using intervals in the computation of the slopes and solving (with an appropriate algorithm - not Gauss elimination!) the resulting linear interval equation, one gets an enclosure

$$N(x) = z - F[z, x]^{-1} F(z)$$

of all slopes, $N(x)$ with $x \in X$.

Slopes IX, ctd.

$N(x)$ is called a Newton operator applied to the box x .
Because of the fixed-point property of $N(x)$, the Newton operator has the important property

Any zero $\hat{x} \in x$ satisfies $\hat{x} \in N(x)$ and hence
 $\hat{x} \in N(x) \cap x$

Thus applying the Newton operator and intersecting with the box never loses a zero (at least with rigorously implemented outward-rounding interval arithmetic)

As a corollary,

$N(x) \cap x = \emptyset \Rightarrow x$ contains no zero

This (and similar) nonexistence result has no analogue in traditional numerical analysis - there is no way to verify nonexistence of solutions by traditional numerical methods (except with symbolic packages)

From the Kakutani Fixed-point theorem
(or the Leray-Schauder P theorem),
one can also deduce

$N(x) \subseteq \text{int } x \Rightarrow x$ contains some zero

This existence result is valid rigorously (in a rigorous implementation)

and is the basis of techniques for computer-assisted proofs in analysis

Slopes X

Interval arithmetic, together with associated existence and nonexistence results (of which we only showed one example each) is

indispensable or reliable global optimization, which is based on pruning the search space by discarding regions which are proved to contain no solution of the problem better than the best already known solution.

The global optimization package BARON is the current state of the art in reliable global optimization and outperforms stochastic (unreliable) search methods on problems up to a few hundred variables. (For larger problems, the search space is often so large that pruning takes too much time, and the extra guarantee of having found the solution is lost if the pruning process is stopped due to a time limit.)

The authors of BARON won in 2006 the Beale-Orchard-Hayer prize, the most prestigious prize awarded by the Mathematical Programming Society.

Slopes XI

The interval Newton method says that

$$\tilde{N}(x) = z - F'(x)^{-1}F(z)$$

has the same properties as the Newton operator discussed here. In fact it predates the work on slopes by almost 20 years. Note that

$$F[z, x] \subseteq F[x, x] = F'(x)$$

so that the formula with the slope is always more accurate.

For example, if $F(x) = x^2$ and $x \in [1, 3]$, $z = 2$

then $F'(x) = 2x = [2, 6]$ has width 4

while $F[z, x] = z + x = [3, 5]$ has only width 2

It can be proved that asymptotically (for narrow boxes) the width of $F[z, x]$ tends to half the width of $F'(x)$, while for larger boxes the gain is often (but not always) much more conspicuous.

These slopes are an essential improvement over derivatives in interval-based applications

Slopes ΔII

To compute interval slopes recursively, one needs enclosures for the slopes of the elementary functions φ . Koler (1997) noticed that for convex and concave functions φ ,

$$\varphi[z, X] = \square \{ \varphi[z, \underline{x}], \varphi[z, \bar{x}] \}, \quad (*)$$

where \square denotes the interval hull. For other functions, a more complicated analysis is needed, but the final result is usually not much more costly than $(*)$.

The real slopes may be computed by difference quotients (any directed rounding to get the outward enclosure property). However, for $z \approx \underline{x}$ or $z \approx \bar{x}$ this may give pessimistic bounds due to cancellation, followed by division by a small number. Therefore it is advisable to intersect the computed result from $(*)$ with $\varphi'(X)$ which is an enclosure (and, as we have seen, overestimates) ^{by a factor of 2} but which does not suffer from stability problems.

Slopes XIII

Higher order derivatives can also be extended to slopes, using

$$f[x_1, x_2, x_3] = f[x_1][x_2, x_3],$$

$$f[x_1, x_2, x_3, x_4] = f[x_1, x_2][x_3, x_4], \text{ etc}$$

These higher order slopes appear as coefficients in the multivariate generalization of Newton's interpolation formula,

$$f(x) = f(x_1) + f[x_1, x_2](x-x_1) + \dots + f[x_1, \dots, x_{k-1}](x-x_1) \dots (x-x_{k-1}) \\ + f[x_1, \dots, x_k, x](x-x_1) \dots (x-x_k).$$

We have used successfully second order slopes for finding good inclusion and exclusion regions for zeros of nonlinear systems. (Schichl & Neumaner 2008)

This generalizes the ~~existence~~ existence and nonexistence results obtainable by (first order) slopes, and sometimes increases the region where an existence or nonexistence decision can be found by order of magnitude.

Higher order slopes can be computed in an AD-like fashion in full analogy to the case of higher derivatives.

Scopes XIV

Some references

My textbook

Introduction to Numerical Analysis

Arnold Neumaier

Cambridge University Press 2001

Covers the standard topics of such a book,
together with forward and reverse automatic differentiation
and interval techniques

(It is probably the only NA textbook covering AD.)

My older text

Interval methods for systems of equations

Cambridge Univ. Press 1990

Covers interval methods in depth but has only
forward slopes (and nothing on the now important
linear relaxation technique discussed below).

The book can still be ordered via print-on-demand
directly from C.U.P. (I still get royalties for about
30 copies each year.)

FMathL I

FMathL (= Formal Mathematical Language) is the working title for a modeling and documentation language for mathematics, to be developed in a 3-year project at Vienna University which just started.

Thus all I say here is about plans, to be realized within the next few years.

The design of FMathL is based on our experience with

- modeling languages
- numerical programming systems
- test problem collections ~~from~~ various problem types
- real applications, and
- wrong starts in the past

FMath L II

The goal of the FMathL project is to

- combine the advantages of LaTeX for writing and viewing mathematics,
- the user-friendliness of mathematical modeling systems such as AMPL for the flexible definition of large-scale numerical analysis,
- the universality of the mathematical language to describe completely arbitrary problems,
- the high-level discipline of the CVX system for solving convex programming problems and enforcing their semantic correctness, and
- the semantic clarity of the \mathbb{Z} notation for the precise specification of concepts and statements.

We believe that this goal is reachable, and that an easy-to-use such system will change the way modeling is done in practice.

F Math L III

Key characteristics of the future language:

- separate completely modeling aspect (declarative) and algorithmic aspect (imperative)
- The model is (almost) the documentation, the latter can be automatically generated from the model in "publishable" quality
- The model is specified close to how it would be communicated informally when describing it in a lecture or paper (except for suppressed details)
- easy exchange of problems on a high level
- very transparent use, fully modular
- user extensible
- public domain
- mathematician friendly
(write as in a textbook...)
- any mathematical problem can be specified (though not necessarily paired to a solver)
They later extensions only affect the interface not the problem structure, and can be done by software contributed by third parties

FMathL IV

Specifications in FMathL

- specify problems in their natural mathematical form, with functions, sets, operators, measures, cases rather than loops, indices, branches
- specify algorithms by their requirements (input conditions, output conditions) rather than a piece of program
- specify strategies for combining algorithms to solve problems from a given category
- document which programs satisfy which algorithmic specifications
 - such that in principle (and perhaps in a followup project) program verification is possible
- From these specifications, automatically create the structure needed to interface with existing programming systems (LAPACK, Matlab, ...) and solvers (for optimization problems, ODEs, PDEs, etc), and with L^AT_EX for documentation.
- As part of the latter, include automatic differentiation capabilities on the highest level:
 - differentiate specifications rather than codes.

Propagator I

What can be propagated successfully through a computational graph?

Traditionally:

- (non)standard data types: single, double, complex, variable precision, interval, polynomial, etc.
- gradients and variants
 $F'(x)u$, $F'(x)^T v$, $f''(x)u$
- Taylor series and higher derivatives
- sparsity patterns
- work counts for all these tasks
- function data type (creating the DAG if no branches are encountered)

In addition, it is possible to propagate a large number of other information ...

Propofahle II

nonstandard

easy to do in forward mode

- ranges (interval data type)
- norm bounds (for error estimation)
↑ quantitative
e.g. $\| \frac{d^k f}{dx^k} \|_{\infty} \leq k! \| f \|_2$ where $(f)_k = \sum f_j g_{k-j}$
etc. in any dimension
- Lipschitz constants $\| F(x) - F(y) \| \leq L \| x - y \|$
- bounds on linear functionals (integrals, high order derivatives, etc.)
using Cauchy's integral theorem
needs ranges on products of discs
or product of boundaries of rectangles $\subseteq \mathbb{C}$

easy to do in backward mode

- error bounds for approximations, e.g. Taylor series
- asymptotic expansions
bounds on unbounded domains
- linear, quadratic, or convex relaxations
also more general structured relaxations such as
separable quadratic, semiseparable, biseparable
(= quadratic + separable) (= sum of bivariables)
- convexity information
Fover: Dr. AMPL ; Yr + Gmpt: CVX
- slopes

- Notes:
- overestimation problems:
on $[-1, 1]$, Lipschitz constant of $x - x^2$ is 2
range of $x - x^2$ is $[-2, 2]$
 - quality / time tradeoff
 - work harder to reduce overestimation
 - extra effort for npar in rounding

Propagata III

The derivation of formulas in the nonstandard cases become fairly easy once one has seen the basic underlying principle. We begin with a traditional backward evaluation to motivate the principle.

Why AD-generated gradients are cheap

Standard evaluation of $f(x)$ at an arbitrary point x gives the values of all intermediate results y_i (where $y_{N+1} = x$) and as final result $y_N = f(x)$.

We can write $f(x)$ in many ways as a function of the intermediate results $y_{1:N}$, for example as

$$f(x) = f_N(y) := y_N.$$

Recursively, we can get a new representation $f_{k-1}(y)$ of $f(x)$ by substituting the defining relation for y_k into $f_k(y)$.

Clearly, $f_k(y)$ depends only on $y_{k:N}$.

In particular, $f_N(y)$ depends only on $y_{N+1} = x$, and indeed, it is the original expression for f as a function of x .

Now suppose that we know the intermediate results y_i^0 for an evaluation at x^0 . Then we can approximate each $f_k(y)$ as a truncated Taylor series $t_k(y)$ centered around y^0 . The substitution process $f_k \rightarrow f_{k-1}$ then clearly amounts to getting $t_{k-1}(y)$ from $t_k(y)$ by substituting the defining relation for y_k into $t_k(y)$ and Taylor expanding the few terms affected by this substitution.

In particular, to get gradients, a linear Taylor expansion suffices. In this case, the only term affected by the substitution is a term $c_k y_k$, $1 \leq k \leq 3$ which, after substitution, can be linearized trivially with $O(1) \leq 3$ operations. Therefore the whole substitution process computes $t_k(y) = f(x^0) + f'(x^0)(x - x^0)$ and hence the gradient in $O(N) \leq 3N$ operations, where N is the number of operations to get $f(x)$ in the first place. Thus gradients are cheap.

To get Hessians, one needs a quadratic Taylor expansion $f + c^T(y - y^0) + \frac{1}{2}(y - y^0)^T G (y - y^0)$. Now the terms to be substituted are of the form $c_k y_k$, $G_{kl} y_l y_k$ ($l < k$) and $\frac{1}{2} G_{kk} y_k^2$. Although typically only a few of these terms are nonzero, one sees that the worst case complexity is significantly worse than $O(N)$.

Proposition IV

Using Newton-polynomials, in place of truncated Taylor polynomials, an identical procedure produces slopes in the linear case, and higher-order slopes in the case of Newton polynomials of higher degree.

Only the details of the expansion change!

Once the generality of the principle is recognized, one can apply it to generate all sorts of desired information. We show how to obtain a linear relaxation of $f(x)$ on a box X , i.e., a linear function $r(x)$ such that

$$r(x) \leq f(x) \text{ for all } x \in X.$$

In an initial forward pass, we compute enclosures y_i for the ranges of all intermediate results y_i . (Actually, this can be improved using constant propagation, see below.)

We construct $r(x)$ as $r_n(y)$, where $r_0(y) = y$, and each $r_{i+1}(y)$ is obtained from $r_i(y)$ by substitution ~~of~~ of y_i and subsequent linear relaxation. As in the case of gradient and slope, the cost is $O(N)$ per step, since we only need to relax the term obtained by substituting into r_i . Thus linear relaxation is cheap.

To be able to program this, we just need to have symbolic linear relaxation for a few expressions, namely

$$\pm (u \circ v) \text{ for } \circ = +, -, *, /, \wedge, \text{ atan2}$$

$$\pm \varphi(u) \text{ for } \varphi = \text{sqrt}, \text{exp}, \text{log}, \text{sin}, \text{cos}, \dots$$

Side conditions for the (nonunique) relaxation of these basic expressions are

either a minimal word corner on the rectangular domain of u, v , resp. the interval for u ,

or a minimal error at a target point u^0, v^0 , obtained by a forward evaluation.

Figuring out an optimal relaxation is now itself a straightforward exercise.

Propagation V

Constraint propagation (CP)

The propagation of ranges can be done in more sophisticated ways when the application is a global problem (for finding a feasible point of a constraint satisfaction problem, or an optimal point of a global optimization problem).

The reason is that in this case the constraints on the results of a function, or an upper bound on the objective given by the best known feasible point, can be used to improve the ranges of intermediate results — though they might be outside the improved ranges for an arbitrary input in the given box, they will be within the range for any feasible input, and only the latter counts.

There be some range information propagated backwards from results to intermediates and ultimately to the original variables. This process is called constraint propagation.

On a computational day, we have always relations of the form

$$z = x \circ y \text{ or } z = \varphi(x)$$

Forward evaluation of ranges requires bounds on z given bounds on x and y , and is standard interval arithmetic, together with direction.

Backward evaluation of ranges requires bounds on x or y given bounds on z (and y or x), and is interval arithmetic on the corresponding inverse function. Complications arise because of the multi-valuedness and because of noninvertibility in a number of cases ($\varphi = \min$, resp. \max), but correct formulas are known in each case.

Ranges that improved significantly are propagated again, also in the forward direction, resulting in a somewhat chaotic propagation scheme that repeatedly traverses the relevant regions of the DAB until improvements become insignificant and a near fixed-point is reached.

Higher consistency techniques based on "sharing" and other tricks are more expensive but give better results on compound expressions.

Linear relaxation

We discussed already how to create linear relaxation by a backward AD-like scheme.

Alternatively, and realized, e.g., in BARON, are linear relaxation techniques that work directly on each relation of the DAB rather than on the original variables only.

BARON (following McCormick (1961)) uses optimal linear inequalities for the standard operators,

$$\text{e.g. } z = xy, \quad x \in [x], \quad y \in [y]$$

$\Rightarrow z \geq \underline{x}y + \underline{y}x - \underline{x}y$ since the difference is $(x-\underline{x})(y-\underline{y}) \geq 0$,
and three other inequalities are obtained from $(x-\underline{x})(y-\bar{y}) \leq 0$
 $(\bar{x}-x)(y-\underline{y}) \leq 0$
 $(\bar{x}-x)(y-\bar{y}) \geq 0$

Accepted analysis for particular compound expressions occurring frequently in chemical expressions (the original target of BARON), such as xy/z , xy/z , etc.

Yields better linear inequalities involving more variables.

Each analysis of a compound expression is a symbolic global optimization problem in few variables and, once solved, can be incorporated into the solver for any constraints containing such a compound subexpression.

However, these symbolic problems get harder with the complexity significantly of the compound.

The family of all ^{linear} relaxations generated is then made the basis of a linear program whose solution reduces selected ranges, frequently with extraordinary success compared to pure CP. The cost of solving these large scale LPs is amortized in BARON by the drastic reduction in the size of the branch-and-bound tree that needs to be considered.

Optimization problems as graphs

$$\begin{array}{l} \max \\ \min \end{array} f(x)$$

$$\text{s.t. } x \in X, F(x) \in Y$$

$$\text{equalities } \hat{=} \in [0, 0]$$

$$\text{inequalities } \hat{=} \in [0, \infty]$$

Introducing variables for slacks and objective

$$\min \pm x_n$$

$$\text{s.t. } x \in X, F(x) = 0$$

with more variables

Enforcing sparse Jacobians and Hessians
by substitution and tearing

$$f(x) = \sum \varphi_i(x_i) + \left(\sum \psi_i(x_i) \right)^2 \quad \begin{array}{l} \text{dense} \\ \text{Hessian} \end{array}$$

$$u = \sum \varphi_i(x_i)$$

$$v = \sum \psi_i(x_i)$$

$$w = u + v [= f(x)]$$

} sparse representation:
Hessian of Lagrangian
is diagonal

Dense rows in Jacobians fill in the least square Hessian term in penalty formulations and interior point codes:

The Hessian of $\frac{1}{2} \|Ax - b\|^2$ involves $A^T A$ which is a dense matrix if A has a dense row: $| \cdot \text{---} = \begin{array}{|c|c|c|c|} \hline \square & \square & \square & \square \\ \hline \end{array}$

Tearing splits an equation

$$(*) \quad \sum f_i(x_{I_i}) = 0 \quad \text{leading to a dense row}$$

$$\text{into} \quad \sum u_i = 0, \quad f_i(x_{I_i}) - u_i = 0$$

which is sparse.

Figuring out the optimal way to tear a constraint into the form (*) is nontrivial

Relation graphs I

Carrying the substitution process to the extreme ends in an optimization problem where every intermediate term is a variable and all constraints are simple way, binary or ternary relations of the form

$$x_i \in \mathcal{X}_i$$

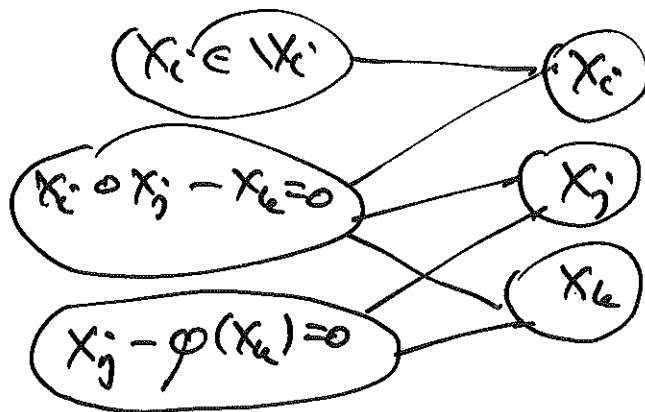
$$\varphi(x_i) - x_j = 0$$

$$x_i \circ x_j - x_k = 0$$

Essentially we get an undirected version of the computational graph.

Substitution and its converse, elimination may then be viewed as transformation of an undirected relation graph, consisting of the union of all edges defined on the bipartite vertex set with relations and variables as vertices, and edges between a variable and a relation if the variable appears in the relation.

E.g.,



By interpreting a variable node as an equality relation between all incoming edges, we can think of all vertices as relations, and all edges as variables.

Relation graphs II

In this interpretation (vertices = relations, edges = variables) all transformations (involving substitution and/or elimination) can be accommodated very naturally.

By focussing on certain types of canonical relations one can bring any optimization problem (with certain restrictions) into a canonical form most useful for applications.

For example, GloptLab (a Matlab global solver under development)

requires a canonical form consisting of quadratic constraints only. Any algebraic problem can be transformed to this form by removing divisions and removing integer powers and roots

e.g., $x/y = z \iff x = yz, y \neq 0$
 $\sqrt{x} = y \iff x = y^2, y \geq 0$

Also many trigonometric problems can be transformed in this way

$$\left. \begin{array}{l} x = r \cos \varphi \\ y = r \sin \varphi \\ r \geq 0, \varphi \in [0, 2\pi] \end{array} \right\} \iff \begin{cases} x^2 + y^2 = r^2 \\ \varphi = \arctan^2(x/y) \end{cases}$$

$$\left. \begin{array}{l} x = \cos 2\varphi \\ y = \sin(\varphi + \alpha) \\ z = \cos \alpha \end{array} \right\} \iff \begin{cases} x = c^2 - s^2, & c^2 + s^2 = 1 \\ y = s w + c z, & w^2 + z^2 = 1 \end{cases}$$

if φ, α are eliminated

But $\varphi + \sin \varphi = 0$ is not reducible to an algebraic problem.

Problem

$$\begin{aligned} \min f &= (4x_1 - x_2x_3)(x_1x_2 + x_3), \\ \text{s.t. } x_1^2 + x_2^2 + x_1x_2 + x_2x_3 + x_2 &= 0, \\ \exp(x_1x_2 + x_2x_3 + x_2 + \sqrt{x_3}) &\in [-1, 1], \\ x_1 \geq 0, \quad x_2 \geq 0, \quad x_3 &\in [-1, 8]. \end{aligned}$$

Code list

$$\begin{aligned} \min f \\ \text{s.t. } f &= v_1v_2 \\ v_1 &= v_3 + v_4 \\ v_2 &= x_4 + x_3 \\ v_3 &= 4x_1 \\ v_4 &= -x_5 \\ v_5 &= x_1^2 \\ v_6 &= x_2^2 \\ v_7 &= \sqrt{x_3} \\ v_8 &= x_6 + v_7 \\ v_9 &= v_5 + v_6 + x_6 \\ v_{10} &= \exp(v_8) \\ x_4 &= x_1x_2 \\ x_5 &= x_2x_3 \\ x_6 &= x_2 + x_4 + x_5 \\ x_1 \geq 0, \quad x_2 \geq 0, \quad x_3 &\in [-1, 8] \\ v_9 = 0, \quad v_{10} &\in [-1, 1]. \end{aligned}$$

relations

min f

s.t. $x_1 = w_{11} = w_{12} = w_{13}$

$$x_2 = w_{21} = w_{22} = w_{23} = w_{24}$$

$$x_3 = w_{31} = w_{32} = w_{33}$$

$$x_4 = w_{41} = w_{42}$$

$$x_5 = w_{51} = w_{52}$$

$$x_6 = w_{61} = w_{62}$$

$$f = v_1 v_2$$

$$v_1 = v_3 + v_4$$

$$v_2 = w_{41} + w_{32}$$

$$v_3 = 4w_{12}$$

$$v_4 = -w_{51}$$

$$v_5 = w_{11}^2$$

$$v_6 = w_{21}^2$$

$$v_7 = \sqrt{w_{33}}$$

$$v_8 = w_{62} + v_7$$

$$v_9 = v_5 + v_6 + w_{61}$$

$$v_{10} = \exp(v_8)$$

$$x_4 = w_{13} w_{22}$$

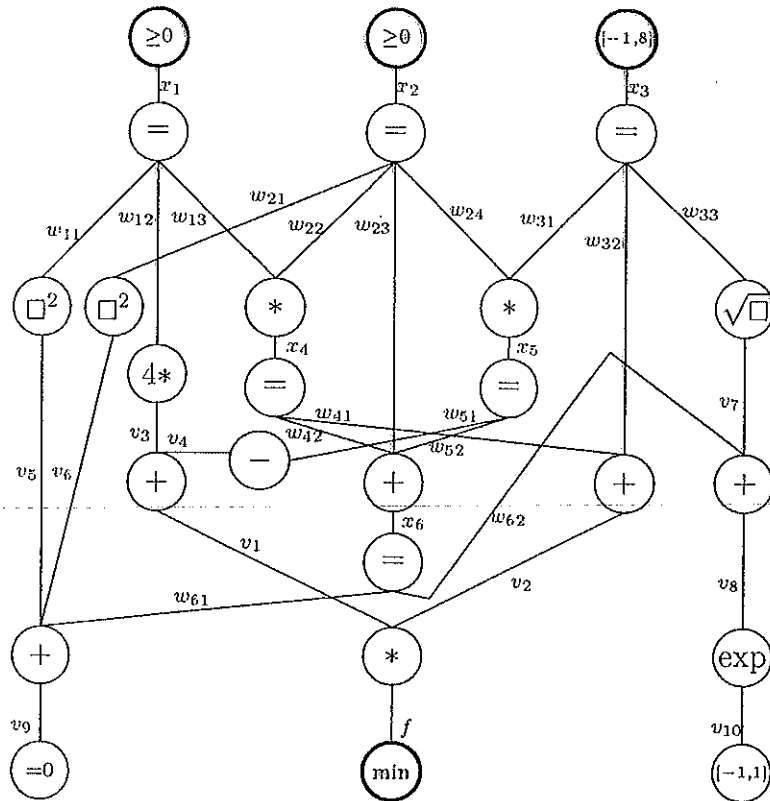
$$x_5 = w_{24} w_{31}$$

$$x_6 = w_{23} + w_{42} + w_{52}$$

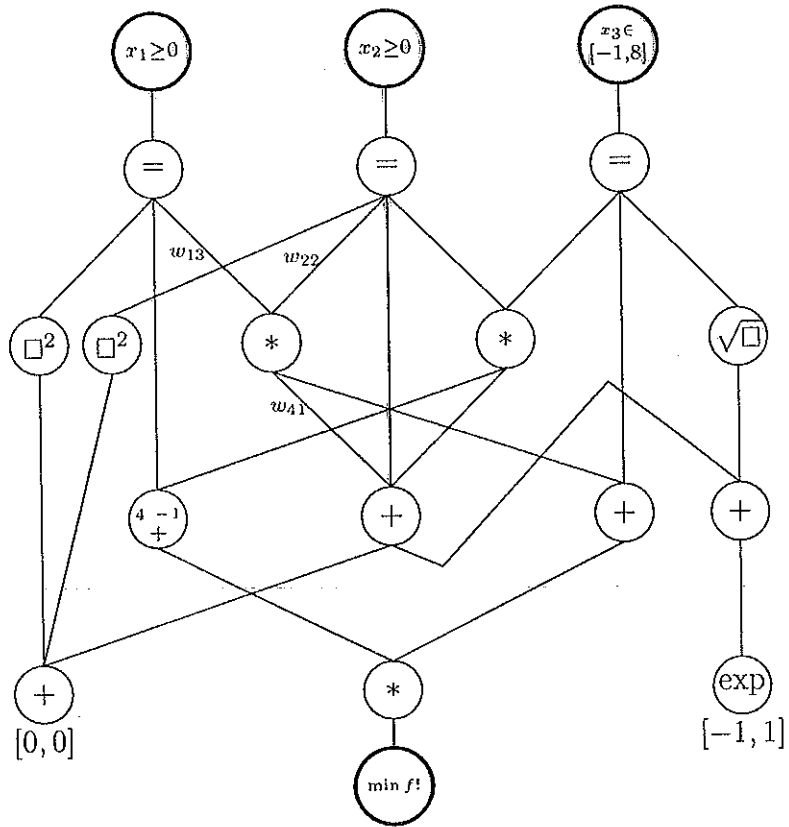
$$x_1 \geq 0, \quad x_2 \geq 0, \quad x_3 \in [-1, 8]$$

$$v_9 = 0, \quad v_{10} \in [-1, 1].$$

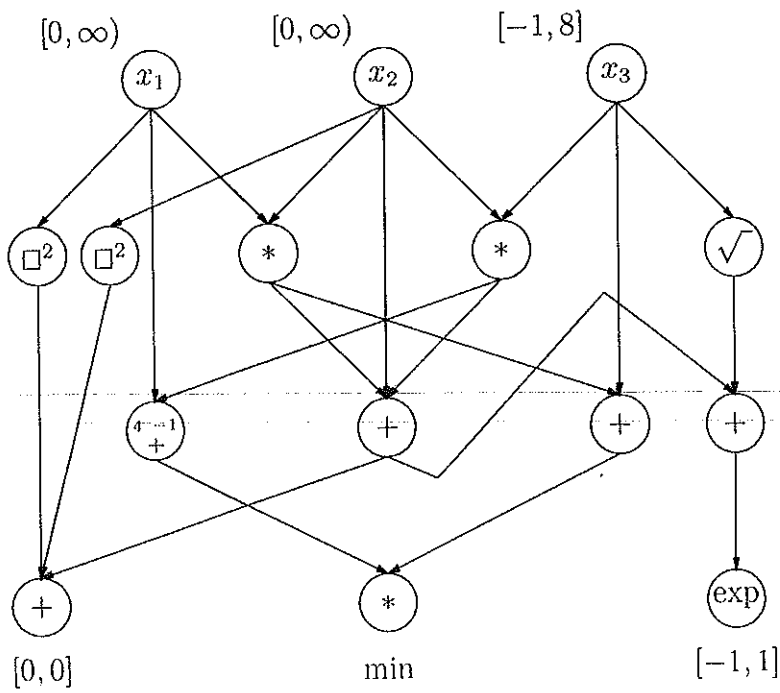
relation graph



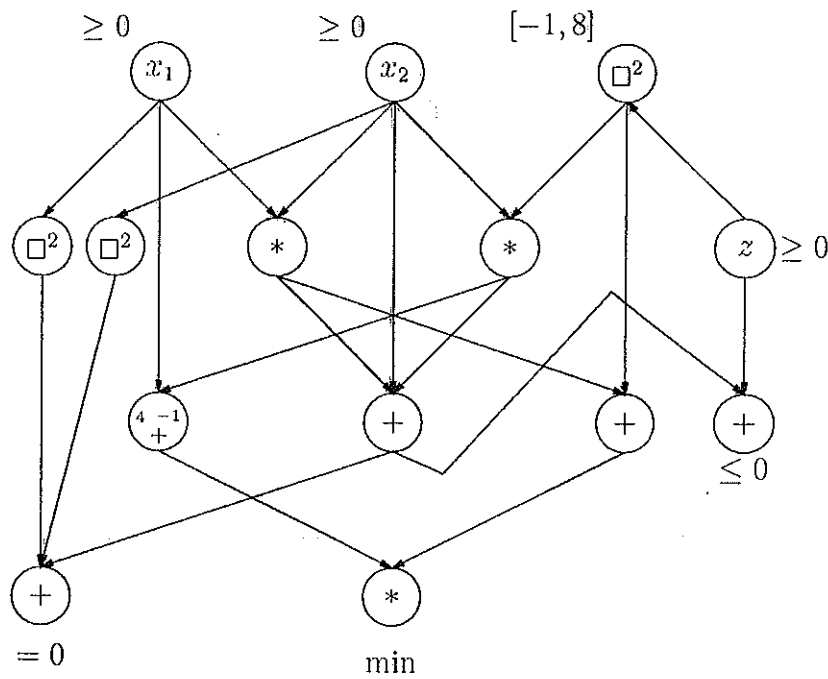
simplex



Oriented as DAG



remove
sqrt
(different DAG)

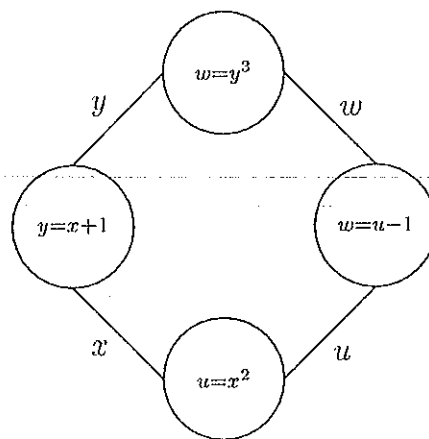
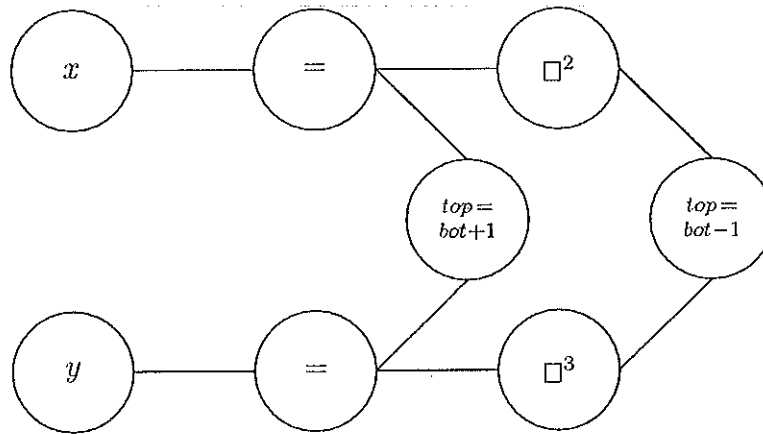


print
problem

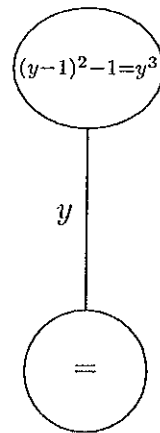
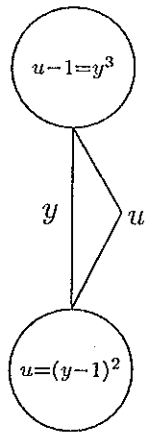
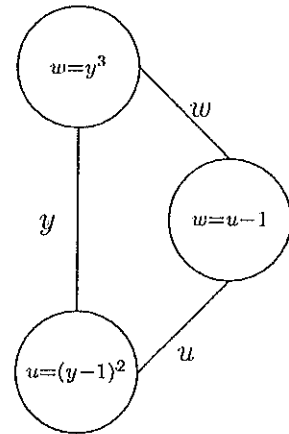
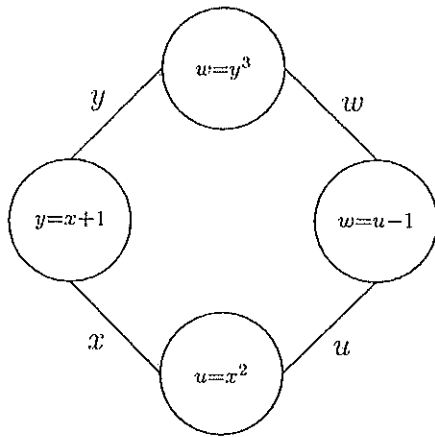
$$\begin{aligned} \min f &= (4x_1 - x_2z^2)(x_1x_2 + z^2), \\ \text{s.t. } x_1^2 + x_2^2 + x_1x_2 + x_2z^2 + x_2 &= 0, \\ x_1x_2 + x_2z^2 + x_2 + z &\leq 0, \\ x_1 \geq 0, \quad x_2 \geq 0, \quad z^2 \in [-1, 8], \quad z &\geq 0. \end{aligned}$$

Graph transformasi

$$y^3 = x^2 - 1$$
$$y = x + 1$$



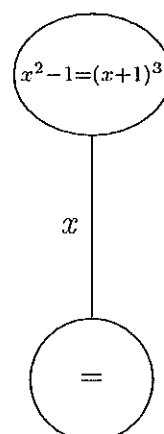
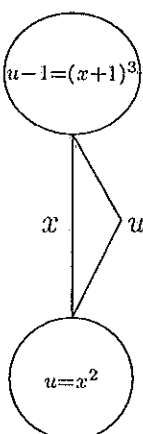
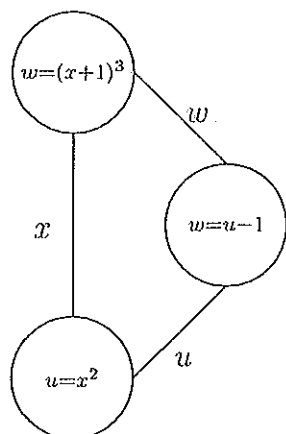
elimination
of x



print
problem

$$y^3 = (y - 1)^2 - 1$$

alternative
elimination
of y



$$x^2 = (x + 1)^3 - 1$$

Relation graphs III

In a relation graph representation in which all intermediate results are made to variables, differentiation is trivial, as is in the more compact representation by separable constraints

$$\frac{1}{2} x^T G x + \sum \varphi_i(x_i) = 0$$

(each φ_i a univariate expression)

What happened to the AD concerns?

The problems of AD on the original problem simply transcribed into problems of sparse matrix analysis!

To see this, note that a DAB evaluation of $z = F(x)$ is nothing else than a solution of the extended sparse triangular nonlinear system $E(x, y) = 0$ for the vector y of intermediate variables, and the projection $z = Py$ which extracts the results from y . By construction, the Jacobian $J = \frac{\partial E}{\partial y}(x, y)$ of this extended system is extremely sparse and lower triangular.

Differentiating, we find $\frac{\partial E}{\partial x} + \frac{\partial E}{\partial y} \frac{\partial y}{\partial x} = 0$, hence

$$F'(x) = \frac{\partial z}{\partial x} = P \frac{\partial y}{\partial x} = - P \underbrace{\left(\frac{\partial E}{\partial y} \right)}_{\substack{\text{forward Jacobian} \\ \text{backward Jacobian} \\ \text{(adjoint)}}} \frac{\partial E}{\partial x} \quad (*)$$

The formula (*) is just the Schur complement of the sparse matrix

$$\begin{pmatrix} \frac{\partial E}{\partial y} & \frac{\partial E}{\partial x} \\ P & 0 \end{pmatrix}$$

Thus AD is just the art of computing certain sparse Schur complements in an optimal way!

Relation graphs IV

Now a Jacobian (or Hessian, which naturally transcribe into Schur complement computations) are never an end in itself.

In ^(local) optimization problems, they are just used to find good search directions or new trial points, by using them in the context of further linear algebra. Thus viewing the Jacobian computation as part of the full problem only may introduce extra flexibility.

This flexibility is naturally provided by the relation graph representation, where the differentiation part and the constraint handling part are seamlessly integrated.

For example, the AD problem of scarcity is naturally transcribed into the graph problem of subdividing out unwound dense pieces, and handling linear pieces within the graph in a special fashion.

On the relation graph level, AD techniques and sparse matrix technology are likely to complement each other in a natural way, and it will be interesting to explore this connection more deeply.

Relation graphs V

The problem transformations enabled by the relation graph framework in a natural and automatic way may even make the difference between a tractable and an intractable problem.

A case demonstrating this nicely is the ^(negative) solution of Wilkinson's conjecture, stating that the growth factor of Gaussian elimination with complete pivoting on an $n \times n$ -matrix is bounded by n . (In fact it grows beyond this conjectured bound)

This is an unconstrained optimization problem, but in this form the problem is intractable for values such as $n=12$. (The number of variables is n^2 .)

By introducing variables for all intermediate elimination matrices, the problem size increases ($O(n^3)$ variables), but it becomes sparse. Although now a constrained problem, the solution for $n=12$ was found easily by Conn, Gould and Torin with their (at that time) new LANCELOT optimization package. The optimal value was slightly above n , thus refuting the conjecture.

Upon analyzing the deeper reason for the success reveals that the expanded constrained problem is much less sensitive to perturbations than the original problem.

In general, scaling and stability issues are easier handled in the sparser formulations.