# Regret Circuits: Composability of Regret Minimizers

Gabriele Farina[1] Christian Kroer[2]
Tuomas Sandholm[1,3,4,5]

[1] Computer Science Department, Carnegie Mellon University
[2] IEOR Department, Columbia University
[3] Strategic Machine, Inc.
[4] Strategy Robot, Inc.
[5] Optimized Markets, Inc.

# Summary of Our Contributions in This Paper

- We introduce a **general methodology for composing regret minimizers**
- Our approach treats the regret minimizers for individual convex sets as **black boxes**
  - Freedom in choosing the best regret minimizer for each individual set
- Several applications, including a significantly **simpler proof of CFR**, the state-of-the-art scalable method for computing Nash equilibrium in large extensive-form games

# Regret Minimizer

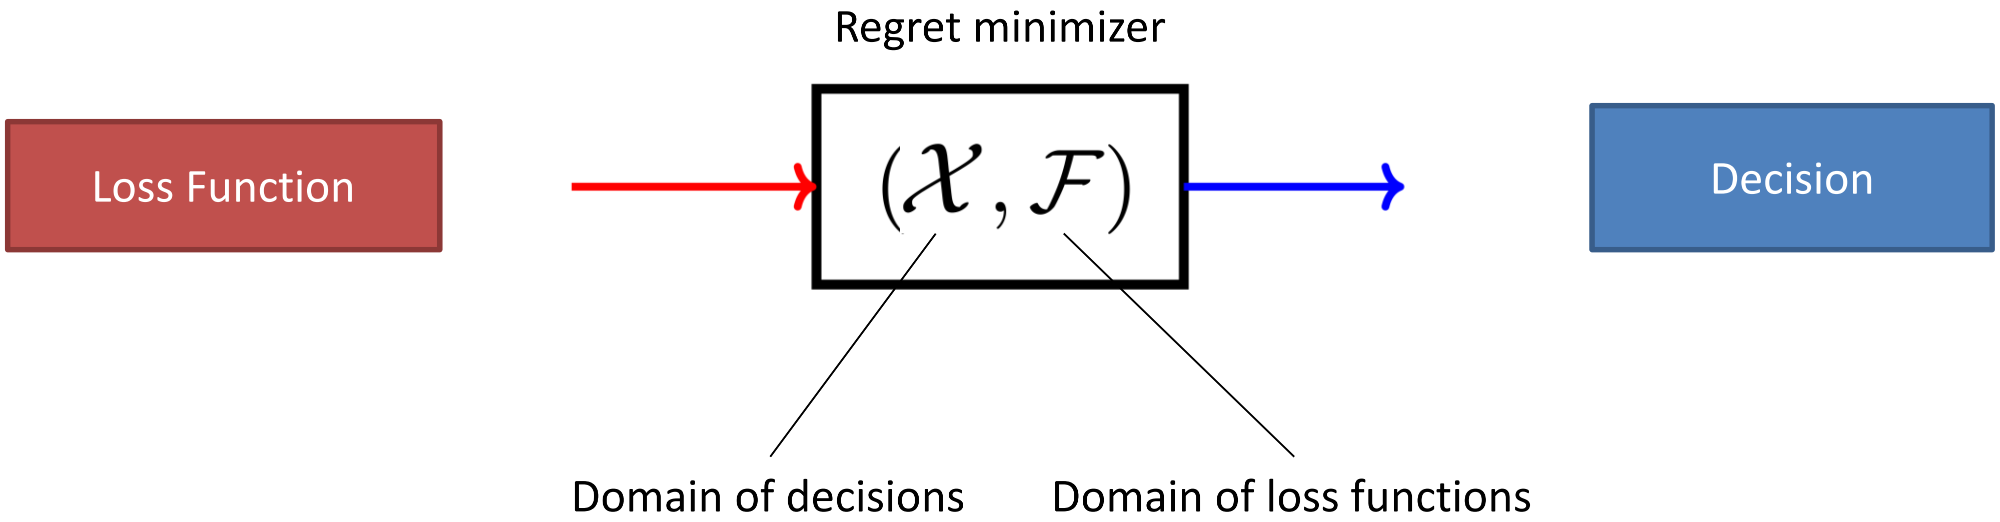

# Cumulative Regret

*"How well do we do against **best, fixed** decision in hindsight?"*

$$R^T := \boxed{\sum_{t=1}^{T} \ell^t(\boldsymbol{x}^t)} - \boxed{\min_{\hat{\boldsymbol{x}} \in X} \left\{ \sum_{t=1}^{T} \ell^t(\hat{\boldsymbol{x}}) \right\}}$$

Loss that was cumulated

Minimum possible cumulative loss

# How to Construct a Regret Minimizer?

- Several “general-purpose” regret minimizers known in the literature:
  - Follow-the-regularized-leader [Shalev-Schwartz and Singer 2007]
  - Online mirror descent
  - Online projected gradient descent [Zinkevich 2003]
  - For simplex domains in particular: regret matching [Hart and Mas-Colell 2000], regret matching+ [Tammellin, Burch, Johanson and Bowling 2000], …
  - …
- Drawbacks of general-purpose methods:
  - Need a notion of **projection** onto the domain of decisions --- this can be expensive in practice!
  - **Monolithic**: they cannot take advantage of the specific (combinatorial) structure of their domain

# Calculus of Regret Minimization

**Idea: can we construct regret minimizers for composite sets by combining regret minimizers for the individual atoms?**

# Easy example: Cartesian product

- How to build a regret minimizer for $X \times Y$ given one for $X$ and one for $Y$?

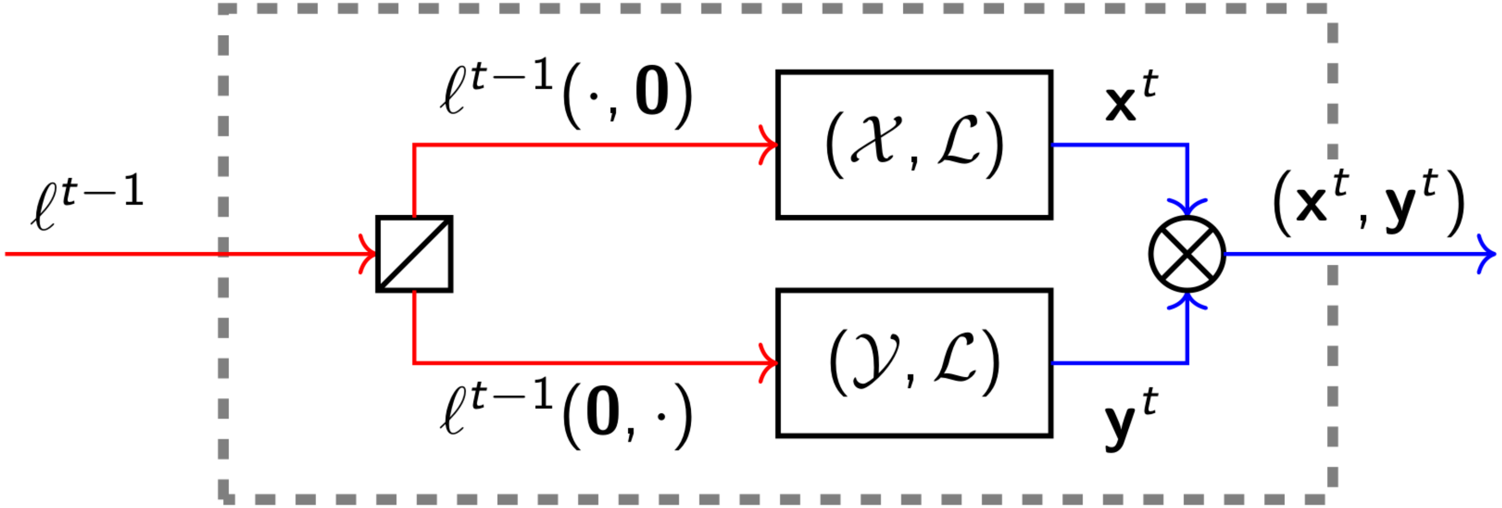


$$R^T = R_X^T + R_Y^T$$

# Harder Example: Convex Hull

- How to build a regret minimizer for the convex hull of $X$ and $Y$ given one for $X$ and one for $Y$?

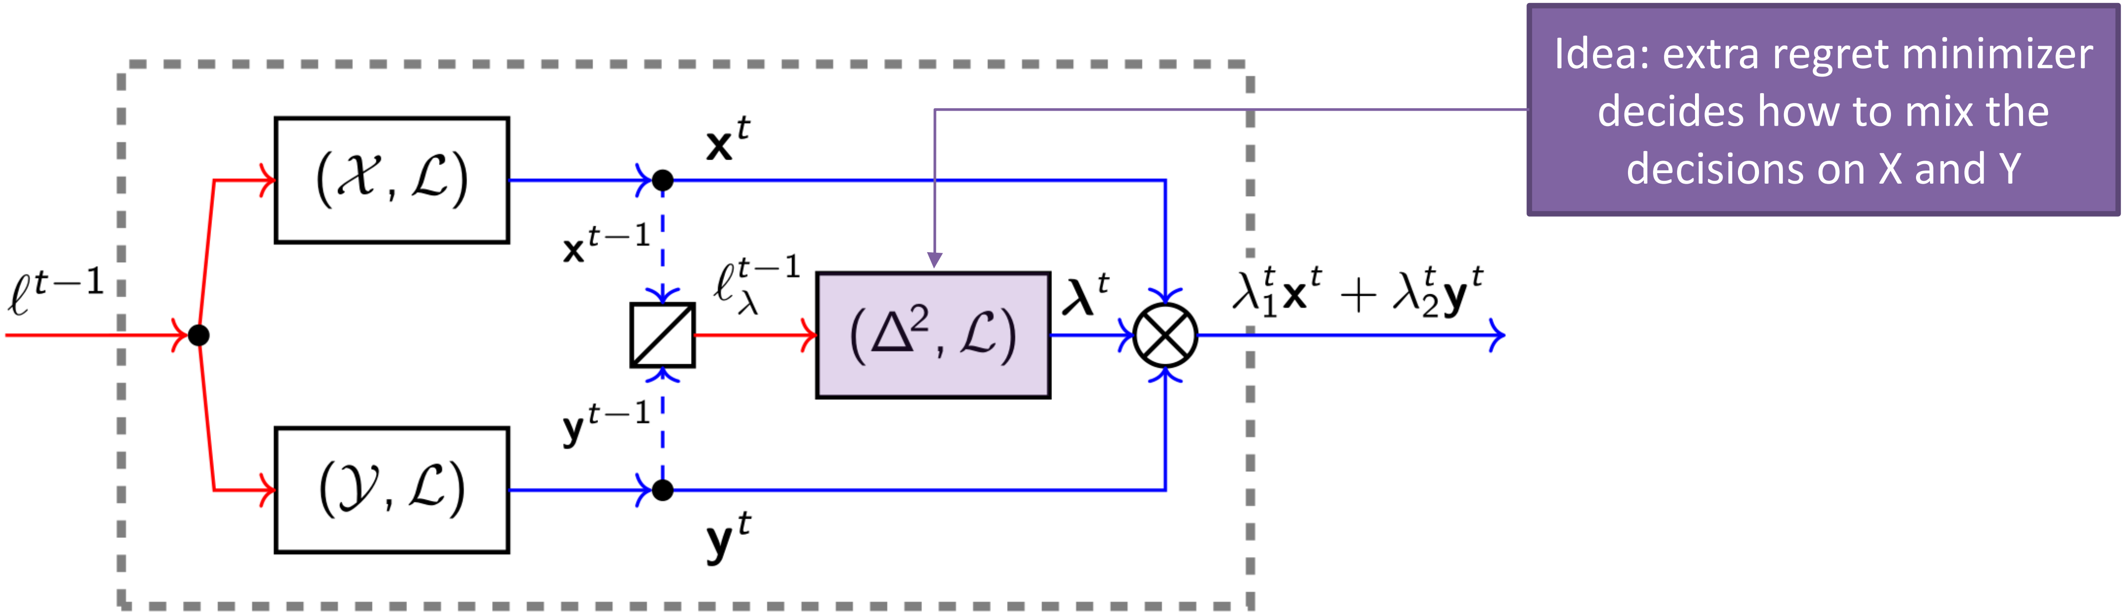


$$R^T \leq R^T_{\Delta^2} + \max\{R^T_X, R^T_Y\}$$

# Intermezzo: Deriving CFR

- Counterfactual regret minimization (CFR) is a family of regret minimizers, specifically tailored for **extensive-form games** [Zinkevich, Bowling, Johanson and Piccione 2007]
- **Practical state of the art for the past 10+ years** in large games
  - One of the key technologies that allowed to solve large Heads-Up Limit and No-Limit Texas Hold'Em [Bowling, Burch, Johanson and Tammelin 2015] [Brown and Sandholm 2017]
- Main insight: break down regret and **minimize it locally** at each decision point in the game
- We can recover the whole, exact CFR algorithm **by simply composing the Cartesian product and convex hull circuits**
  - This also includes newer variants such as CFR+ [Tammellin, Burch, Johanson and Bowling 2015] and DCFR [Brown and Sandholm 2019]

# Intermezzo: Deriving CFR

- Idea: the space of strategies of a player can be expressed inductively by using convex hulls and Cartesian products

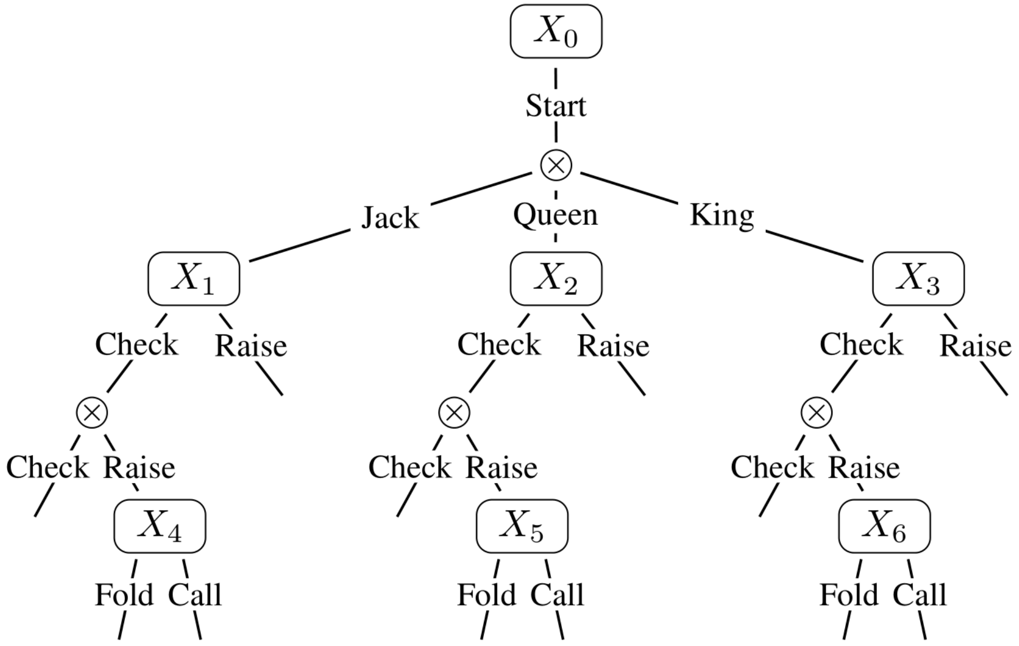


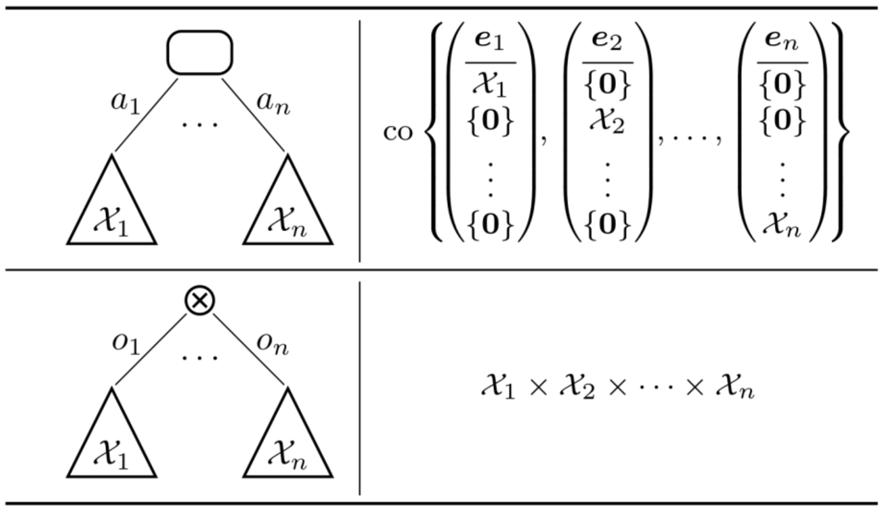

# Calculus of Regret Minimization (cont'd)

- What about **intersections and constraint satisfaction**? We show two different circuits:
  - Approximate circuit using Lagrangian relaxation
  - Exact circuit using (generalized) projections

# Constraint Satisfaction (Lagrangian Relaxation)

- How to build a regret minimizer for $X \cap \{\boldsymbol{x}: g(\boldsymbol{x}) \leq 0\}$ given one for $X$?

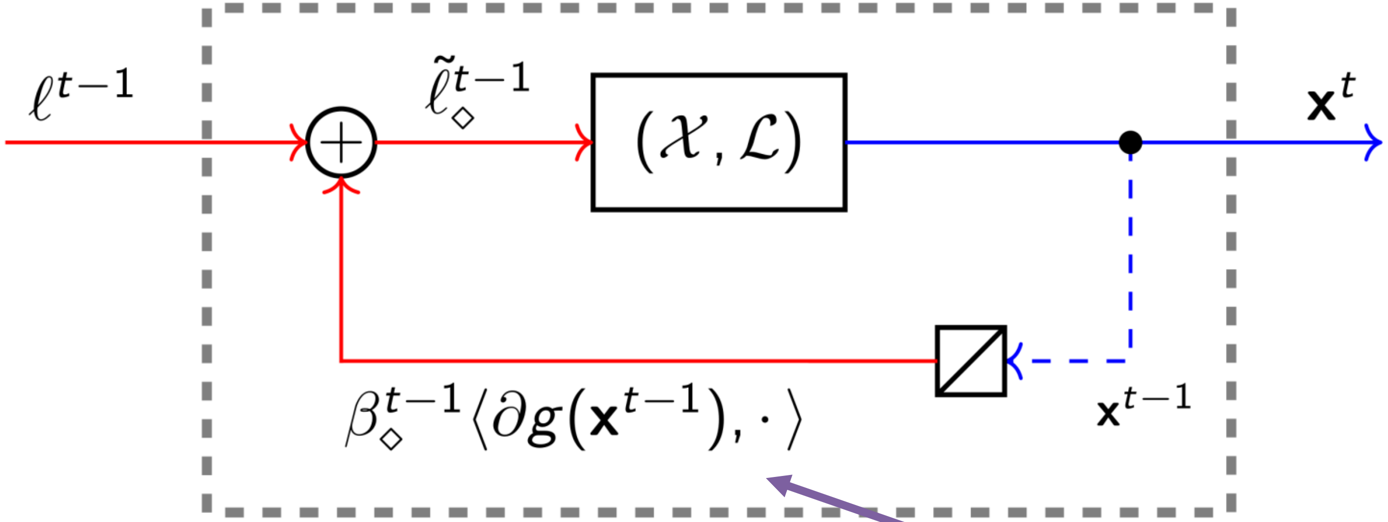


$$\beta_{\diamond}^{t} = \begin{cases} \beta^t & \text{if } g(\mathbf{x}^t) > 0 \\ 0 & \text{otherwise} \end{cases}$$

**Penalization term**!
How feasible was the last recommendation?

# Intersection Circuit

- Want feasibility? Project onto the feasible set!
- Generalized projections (proximal operators) can be used as well

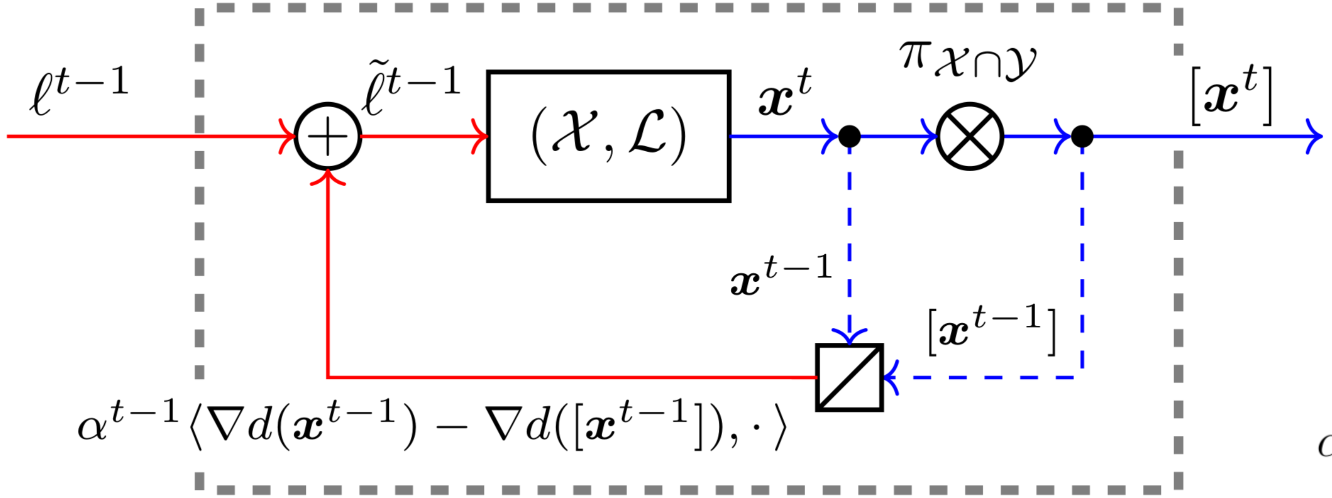


**Penalization term:**

$$\alpha^t = \begin{cases} 0 & \text{if } \boldsymbol{x}^t \in \mathcal{X} \cap \mathcal{Y} \\ \max\left\{0, \dfrac{\ell^t([\boldsymbol{x}^t] - \boldsymbol{x}^t)}{\mu \|[\boldsymbol{x}^t] - \boldsymbol{x}^t\|^2}\right\} & \text{otherwise} \end{cases}$$

- Takeaway: we can always turn an infeasible regret minimizer into a feasible one by projecting onto the feasible set, **outside the loop**!

# Second Intermezzo: CFR with Strategy Constraints

- The recent Constrained CFR algorithm [Davis, Waugh and Bowling, 2019] can be constructed as a special example via our framework, by using the Lagrangian relaxation circuit
- Our exact (feasible) intersection construction leads to a new algorithm for the same problem as well
- Tradeoff between feasibility and computational cost
  – Projections are expensive in general
  – Feasibility might be crucial depending on the application

# Another Application: Optimistic/Predictive Regret Minimization

- A related calculus of regret minimization can be designed for **optimistic** regret minimization
- Optimistic regret minimization breaks the learning-theoretic barrier $O(T^{-1/2})$ on the convergence rate of regret-based approaches
- We use our calculus to prove that under certain hypotheses CFR can be modified to have a convergence rate of $O(T^{-3/4})$ to Nash equilibrium, instead of $O(T^{-1/2})$ as in the original (non-optimistic) version [Farina, Kroer, Brown and Sandholm, 2019]

# Another Application: Extensive-Form Perfect Equilibrium

- We give the first **efficient** regret minimizer for computing **extensive-form correlated equilibrium** in large two-player games [Farina, Ling, Fang and Sandholm, under review]
  - Solution concept in which the game is augmented with a mediator that can recommend behavior but not enforce it --- recommended behavior must be incentive compatible
  - Can lead to very interesting/nonviolent behavior in extensive-form games such as Battleship
- Significantly more challenging than designing one for the Nash equilibrium counterpart, as the constraints that define the space of correlated strategies lack the hierarchical structure and might even form cycles
  - We unroll this space without using intersection!

# Another Application: Extensive-Form Perfect Equilibrium

- We use a different regret circuit, for a convexity-preserving operation that we call **scaled extension**

$$\mathcal{X} \overset{h}{\triangleleft} \mathcal{Y} := \{(\boldsymbol{x}, \boldsymbol{y}) : \boldsymbol{x} \in \mathcal{X},\ \boldsymbol{y} \in h(\boldsymbol{x})\mathcal{Y}\}.$$

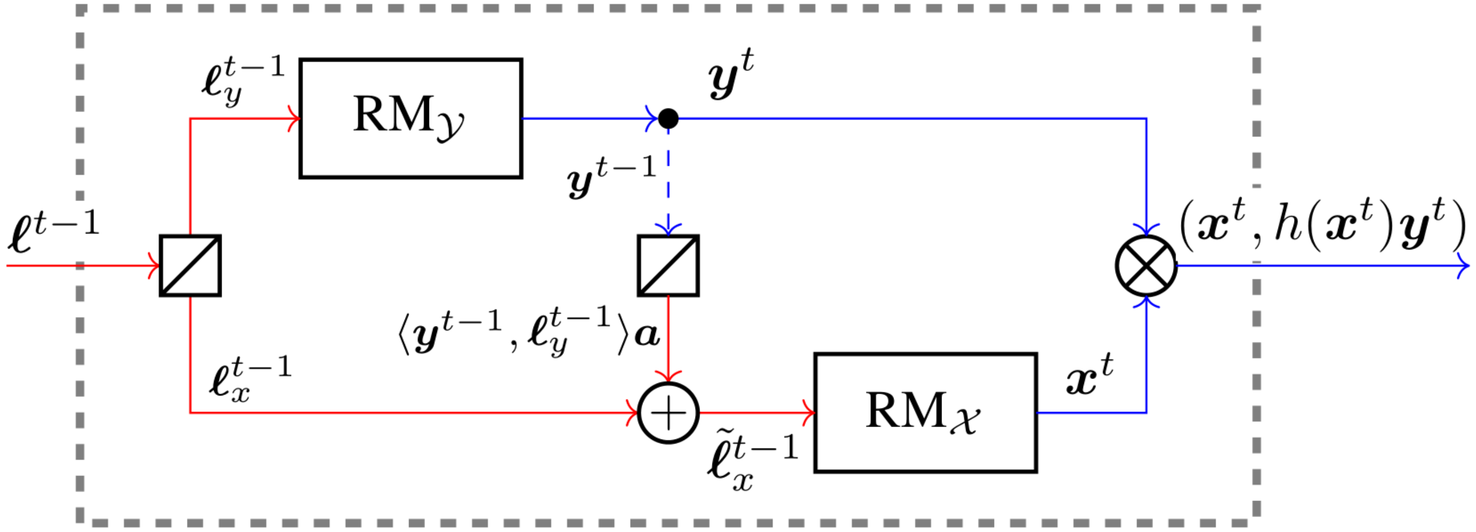

# Conclusions

- We initiated the study of a calculus of regret minimizers
  - Regret minimizers are combined as black boxes. Freedom to chose the best algorithm for each set that is being composed
  - In the paper we show regret circuits for several convexity-preserving operations (convex hull, Cartesian product, affine transformations, intersections, Minkowski sums, …)
- Our framework has many applications:
  - CFR, the state-of-the-art algorithm for Nash equilibrium in large games, falls out almost trivially as a repeated application of only two circuits
  - Improves on the recent ‘CFR with strategy constraints’ algorithm
  - Leads to the first CFR variant to beat the $O(T^{-1/2})$ convergence rate when computing Nash equilibria
  - Gives the first efficient regret minimizer for extensive-form correlated equilibrium in large games

# Future research

- Full generality over the class of functions
  - Most circuits assume linear losses
  - What about general convex losses?
- Deriving a full calculus of *optimistic/predictive* regret minimization
  - So far: only convex hulls and Cartesian products
- Improving on the intersection construction in special cases
- More circuits for specialized applications

Poster: **Pacific Ballroom #150** 06:30 - 09:00 pm