Ngu (Nathan) Dang

Research Statement

✉ ndang@bu.edu | 🔗 nathandang97.github.io | linkedin.com/in/ngu-dang

Overview

My research is in *computational complexity*, focusing on **circuit complexity** and its connection to **meta-complexity**, which investigates the resources needed to evaluate or certify complexity properties (such as circuit size) of explicit objects. In particular, I study how popular and fundamental techniques (such as Gate Elimination) can be used to characterize optimal Boolean circuitsfan-in/fan-out, wiring patterns, and compositional "shapes" and also analyze why these standard methods stall for some particular cases. A recurring theme is to take familiar and fundamental technique and reinforce them with principled invariants and decomposition frameworks, so we either (1) push known lower bounds beyond classic barriers or (2) explain where a technique fails, thereby guiding new approaches.

Beyond these structural questions, I explore how meta-complexity questions (e.g., the algorithmic detectability of small circuits) can leverage structural characterizations of explicit functions. My long-term goal is to build a toolkit that connects circuit shape, provable lower bounds, and meta-complexity hardness. I am also interested in connecting this line of work to cryptographic heuristics and assumptions (e.g. existence of one-way functions) and to other computational models such. More broadly, I hope to connect these structural and meta-complexity questions with neighboring areas of theoretical computer science, and I am open to learning new topics and techniques whenever they offer a route to such connections.

1. Research Highlights

My recent work has two threads. First, I develop structure theorems for optimal circuits of explicit functions (e.g. the **XOR**-function), pinning down not just the sizes of best circuits, but what they look like. Second, I study small additive extensions (or *simple extensions* for short) which asks: starting from an optimal circuit for a base Boolean function f, can we compute a slightly larger target function g by adding only "a few" new gates to an optimal circuit for f? This viewpoint connects circuit lower bounds and optimal structures to meta-complexity problems such as the well-known Minimum Circuit Size Problem (**MCSP**).

1.1. Characterizing Optimal Boolean Circuits

Problem Setting. For a Boolean function f and a complexity measure μ, we ask whether characterizing μ-optimal circuits computing f is feasible. When μ is measured under the DeMorgan basis, the answer is trivial for $f = \mathbf{OR}$; it is historically delicate with intricate case analyses for $f = \mathbf{XOR}$ when NOT-gates are counted towards μ (Kombarov, 2011); and it remains open for the Multiplexer $f = \mathbf{MUX}$.

My Contribution. In a joint work with Marco Carmosino (IBM) and Tim Jackman (Boston University), we show that even when NOT-gates are free, any minimal DeMorgan circuit for $\mathbf{XOR}_n$

is essentially a binary tree of $(n-1)$ $\mathbf{XOR}_2$-sub-circuits. This nails down a canonical structure for optimal **XOR**-circuits that matches its classical $3(n-1)$ size bound (Schnorr, 1974). Then, I push forward on **MUX**, a practical candidate whose lower bound (Paul, 1975) and upper bound (Klein & Patterson, 1980) are asymptotically close. In particular, I sharpen its classic upper bound construction with cleaner, level-wise accounting and explain why Gate Elimination alone hits an obstacle for improving its current known circuit lower bound. This analysis motivates a potential alternative route using static cover/"fusion" ideas as a lower bound technique known as *Fusion Method* (Wigderson, 1993; Cavalar & Oliveira, 2025), should Gate Elimination alone is not enough to close the gap between the known lower bound and upper bound.

Why it matters. A concrete circuit shape theorem lifts "existential optimality" to structural optimality which is useful both for excluding whole classes of non-optimal constructions and for algorithms and reductions in meta-complexity that require reasoning beyond mere circuit size.

1.2. Simple Circuit Extensions in Meta-Complexity

Problem Setting. Suppose we start from an optimal circuit for a base function f. We ask whether we can obtain a *simple extension* of f to a target g by adding only "a few" gates to an optimal circuit computing f. When the answer is *yes*, then if optimal circuits for f are small, so are those of g. This (additive) simple extension problem reduces to **MCSP** which, in turns, yields a useful lens in meta-complexity related problems: if recognizing such extensions is easy for certain f, that closes a possible route to proving hardness for **MCSP**; if it is hard, that route stays open. This framework was used to prove hardness of **MCSP** for partial functions under the Exponential Time Hypothesis (ETH) (Ilango, 2020).

My Contribution. In a joint work with Marco Carmosino (IBM) and Tim Jackman (Boston University), we formalize the simple extension problem and show that for **XOR**, it is decidable in polynomial time. In particular, we develop a fixed-parameter algorithm for finding simple extension of some Boolean function f and it is tractable whenever (1) f has linear circuit complexity, (2) optimal circuits computing f have constant fan-out, and (3) there are polynomially many canonical optimal circuits for f with respect to the length of its truth table 2^n. Our result on **XOR**'s rigid "binary-tree" shape satisfies all three conditions, which is why the decision problem when $f = \mathbf{XOR}$ is tractable, and in turns, rules out **XOR** as a candidate for extending known hardness routes to total **MCSP** via Ilango's framework.

Why it matters. Our result gives clear conditions that any potential candidate for proving ETH-hardness of **MCSP** must defeat. This points to **MUX** as a more promising candidate, because its known best circuit construction includes components with *non-constant* fan-out. Thus, if this construction is indeed optimal, our fixed-parameter algorithm for finding simple extension of $f = \mathbf{MUX}$ will run in super-polynomial time.

2. My Research Plans

My aim is to motivate the systematic study of optimal Boolean circuit structure, because such characterizations already yield strong consequences in meta-complexity, where many questions remain open. To that end, I outline my near-term and long-term plans below.

Upcoming Plans. In the near term, my research will proceed along three closely related directions. First, I aim to characterize optimal circuits for other explicit functions that may exhibit a similar “tree-shaped” structure to **XOR**, with the goal of ultimately ruling out these functions as viable candidates for proving hardness of **MCSP**. Second, I plan to tighten the lower bound for **MUX**, ideally using both Gate Elimination and the Fusion Method, in order to gain deeper insight into the optimal circuit structure of **MUX** and to better understand the respective strengths and limitations of these techniques. Third, I will study the simple extension problem for **MUX**, analyzing whether our fixed-parameter algorithm succeeds or fails in this setting and what this reveals about the complexity of **MCSP**. Together, these steps reinforce the arc of my current research: from circuit structure, to new lower-bound techniques, to applications in lower bounds and meta-complexity.

Long-term Plans. Beyond **MUX**, I aim to extend these results to other computational models (e.g., branching programs) and to explore connections to meta-complexity problems related to **MCSP**, such as *time-bounded Kolmogorov complexity*, which is tightly connected to the existence of one-way functions (Liu & Pass, 2020). In other words, my long-term vision is to establish conditions under which structural optimality yields strong complexity-theoretic consequences, e.g. lower bounds and hardness across multiple models. With a deeper understanding of **MUX**, I also aim to develop a toolkit for characterizing optimal circuit structures for Boolean functions in general, integrating both Gate Elimination and the Fusion Method. In pursuing these goals, I am open to situate my work within a broader landscape of theoretical computer science and to broaden my expertise by learning new topics and neighboring areas, e.g. proof complexity, learning theory, and methods/techniques that arise from these connections.