# Communication-Optimal Loop Nests 

Nicholas Knight

New York University<br>Dept. of Mathematics and<br>Center for Data Science<br>nknight@nyu.edu

SIAM Annual Meeting July 13, 2017

Joint work with:

- Michael Christ, UC-Berkeley, Dept. of Mathematics
- James Demmel, UC-Berkeley, Depts. of Mathematics and EECS
- Thomas Scanlon, UC-Berkeley, Dept. of Mathematics
- Katherine Yelick, Lawrence Berkeley Natl. Lab. and UC-Berkeley, Dept. of EECS


## Modeling Communication Cost

## Matrix Multiplication <br> (tiled vs. untiled implementation)



## Modeling Communication Cost

## Matrix Multiplication <br> (tiled vs. untiled implementation)



Observation: \# operations is generally unreliable for predicting runtime - neglects the cost of communication.

## Modeling Communication Cost



Observation: \# operations is generally unreliable for predicting runtime - neglects the cost of communication.

Model Machine:


## Communication Cost $\mathcal{C}$

$\mathcal{C}=\#$ loads/stores

## Motivating Application: Matrix Multiplication

Matrix multiplication algorithm:

$$
\begin{aligned}
& \text { for } i=1: N \\
& \text { for } j=1: N \\
& \quad \text { for } k=1: N \\
& \quad \mathbf{C}_{i, j} \leftarrow \mathbf{C}_{i, j}+\mathbf{A}_{i, k} * \mathbf{B}_{k, j}
\end{aligned}
$$

## Motivating Application: Matrix Multiplication

Matrix multiplication algorithm:

$$
\begin{aligned}
& \text { for } i=1: N \\
& \text { for } j=1: N \\
& \quad \text { for } k=1: N \\
& \quad \mathbf{C}_{i, j} \leftarrow \mathbf{C}_{i, j}+\mathbf{A}_{i, k} * \mathbf{B}_{k, j}
\end{aligned}
$$

Geometric interpretation:

- Operations are points $(i, j, k)$ in Euclidean space;
- Operands are projections $(i, k),(k, j),(i, j)$ onto the coordinate planes.


## Matrix Multiplication: Untiled Implementation

Untiled implementation:

$$
\text { for } i=1: N
$$

$$
\text { for } j=1: N
$$

$$
\text { for } k=1: N
$$

$$
\text { load } \mathbf{A}_{i, k}, \mathbf{B}_{k, j}, \mathbf{C}_{i, j}
$$

$$
\mathbf{C}_{i, j} \leftarrow \mathbf{C}_{i, j}+\mathbf{A}_{i, k} * \mathbf{B}_{k, j}
$$

$$
\text { store } \mathbf{C}_{i, j}
$$

$\mathcal{C} \approx N^{3} \quad$ (no data reuse)


## Matrix Multiplication: Tiled Implementation

Tiled implementation:

$$
\begin{aligned}
& \text { for } \hat{\imath}=1: b: N \\
& \text { for } \hat{\jmath}=1: b: N \\
& \text { for } \hat{k}=1: b: N \\
& \quad \text { load blocks } \mathbf{A}_{\hat{\imath}, \hat{k}}, \mathbf{B}_{\hat{k}, \hat{\jmath}}, \mathbf{C}_{\hat{\imath}, \hat{\jmath}} \\
& \quad \text { for } i=\hat{\imath}: \hat{\imath}+b-1 \\
& \quad \text { for } j=\hat{\jmath}: \hat{\jmath}+b-1 \\
& \quad \text { for } k=\hat{k}: \hat{k}+b-1 \\
& \mathbf{C}_{i, j} \leftarrow \mathbf{C}_{i, j}+\mathbf{A}_{i, k} * \mathbf{B}_{k, j} \\
& \text { store block } \mathbf{C}_{\hat{\imath}, \hat{\jmath}}
\end{aligned}
$$



$$
\mathcal{C} \approx \frac{N^{3}}{b} \quad(b \text {-fold data reuse }) \quad(b=1: \text { untiled impl. })
$$

## Matrix Multiplication: Tiling Attains Lower Bound

Picking tiling parameter $b$ to minimize communication cost $\mathcal{C}$

- $b$-fold reuse means we want to maximize $b$.
- A-, B-, C-blocks must fit in fast memory $\left(3 b^{2} \leq M\right)$.

$$
\text { Picking } b \approx M^{1 / 2} \quad \Rightarrow \quad \mathcal{C} \approx \frac{N^{3}}{M^{1 / 2}}
$$

## Matrix Multiplication: Tiling Attains Lower Bound

Picking tiling parameter $b$ to minimize communication cost $\mathcal{C}$

- $b$-fold reuse means we want to maximize $b$.
- A-, B-, C-blocks must fit in fast memory $\left(3 b^{2} \leq M\right)$.

$$
\text { Picking } b \approx M^{1 / 2} \quad \Rightarrow \quad \mathcal{C} \approx \frac{N^{3}}{M^{1 / 2}}
$$

## Communication Lower Bound (Hong-Kung, 1981)

In any implementation of matrix multiplication, $\mathcal{C} \succeq \frac{N^{3}}{M^{1 / 2}}$.

## Matrix Multiplication: Tiling Attains Lower Bound

Picking tiling parameter $b$ to minimize communication cost $\mathcal{C}$

- $b$-fold reuse means we want to maximize $b$.
- A-, B-, C-blocks must fit in fast memory $\left(3 b^{2} \leq M\right)$.

$$
\text { Picking } b \approx M^{1 / 2} \quad \Rightarrow \quad \mathcal{C} \approx \frac{N^{3}}{M^{1 / 2}}
$$

## Communication Lower Bound (Hong-Kung, 1981)

In any implementation of matrix multiplication, $\mathcal{C} \succeq \frac{N^{3}}{M^{1 / 2}}$.

- Iteration space tiling: heuristic to reduce $\mathcal{C}$ in loop nests.
- This loop nest: tiling with $b$-by- $b$-by- $b$ cubes minimizes $\mathcal{C}$.


## Lower Bound Ingredient \#1: Loomis-Whitney Inequality



Array subscripts are linear maps from $\mathbb{Z}^{3}$ to $\mathbb{Z}^{2}$ :

$$
\begin{aligned}
& \phi_{\mathbf{A}}:(i, j, k) \mapsto(i, k), \\
& \phi_{\mathbf{B}}:(i, j, k) \mapsto(k, j), \\
& \phi_{\mathbf{C}}:(i, j, k) \mapsto(i, j) .
\end{aligned}
$$

## Lower Bound Ingredient \#1: Loomis-Whitney Inequality



## Loomis-Whitney Inequality

For any $S \subseteq \mathbb{Z}^{3}$,
$|S| \leq\left|\phi_{\mathbf{A}}(S)\right|^{1 / 2} \cdot\left|\phi_{\mathbf{B}}(S)\right|^{1 / 2} \cdot\left|\phi_{\mathbf{C}}(S)\right|^{1 / 2}$.

Array subscripts are linear maps from $\mathbb{Z}^{3}$ to $\mathbb{Z}^{2}$ :

$$
\begin{aligned}
& \phi_{\mathbf{A}}:(i, j, k) \mapsto(i, k), \\
& \phi_{\mathbf{B}}:(i, j, k) \mapsto(k, j), \\
& \phi_{\mathbf{C}}:(i, j, k) \mapsto(i, j) .
\end{aligned}
$$

## Lower Bound Ingredient \#1: Loomis-Whitney Inequality



## Loomis-Whitney Inequality

For any $S \subseteq \mathbb{Z}^{3}$,
$|S| \leq\left|\phi_{\mathbf{A}}(S)\right|^{1 / 2} \cdot\left|\phi_{\mathbf{B}}(S)\right|^{1 / 2} \cdot\left|\phi_{\mathbf{C}}(S)\right|^{1 / 2}$.

The memory footprint of iterations $S$ is

Array subscripts are linear maps from $\mathbb{Z}^{3}$ to $\mathbb{Z}^{2}$ :

$$
\begin{array}{ll}
\phi_{\mathbf{A}}:(i, j, k) \mapsto(i, k), & \text { Loomis-Whitney gives the lower bound, } \\
\phi_{\mathbf{B}}:(i, j, k) \mapsto(k, j), & \mu(S) \geq|S|^{2 / 3} . \\
\phi_{\mathbf{C}}:(i, j, k) \mapsto(i, j) . &
\end{array}
$$

## Lower Bound Ingredient \#2: Segmentation

$\quad:$
operation
load
store
operation
load
load
operation
store
store
load
operation
store
load
segment $i$
iterations $S_{i}$
$\left|S_{i}\right|=2$
$\vdots$

## Lower Bound Ingredient \#2: Segmentation

For any segment $i$,
operation
load
store
operation load load
operation store store load
operation store load
-
-
-

$$
\begin{aligned}
\mathcal{C}_{i} & \geq \mu\left(S_{i}\right)-2 M \\
& \geq\left|S_{i}\right|^{2 / 3}-2 M .
\end{aligned}
$$

$$
\begin{gathered}
\text { s } \\
{\left[\begin{array}{l}
\text { it } \\
\text { in }
\end{array}\right.}
\end{gathered}
$$

iterations $S_{i}$ $\left|S_{i}\right|=2$
segment $i$

$$
\left|S_{i}\right|=2
$$

## Lower Bound Ingredient \#2: Segmentation

- 
- 

operation
load
store
operation load
load
operation store store load
operation store load
-
.
-

For any segment $i$,

$$
\begin{aligned}
\mathcal{C}_{i} & \geq \mu\left(S_{i}\right)-2 M \\
& \geq\left|S_{i}\right|^{2 / 3}-2 M .
\end{aligned}
$$

Summing over all segments,

$$
\mathcal{C} \geq \sum_{i}\left|S_{i}\right|^{2 / 3}-2 M
$$

$\left[\begin{array}{l}\text { segment } i \\ \text { iterations } S_{i} \\ \left|S_{i}\right|=2\end{array}\right.$
$\left[\begin{array}{l}\text { segment } i \\ \text { iterations } S_{i} \\ \left|S_{i}\right|=2\end{array}\right.$
$\left[\begin{array}{l}\text { segment } i \\ \text { iterations } S_{i} \\ \left|S_{i}\right|=2\end{array}\right.$

## Lower Bound Ingredient \#2: Segmentation

;
$\square$
operation
load
store
operation load
load
operation store store load
operation
store
load
-
■
-

For any segment $i$,

$$
\begin{aligned}
\mathcal{C}_{i} & \geq \mu\left(S_{i}\right)-2 M \\
& \geq\left|S_{i}\right|^{2 / 3}-2 M .
\end{aligned}
$$

Summing over all segments,

$$
\mathcal{C} \geq \sum_{i}\left|S_{i}\right|^{2 / 3}-2 M
$$

If $\left|S_{i}\right| \approx M^{3 / 2}$ for all segments $i$,

$$
\mathcal{C} \succeq \frac{N^{3}}{M^{3 / 2}} \cdot M=\frac{N^{3}}{M^{1 / 2}}
$$

the stated lower bound.

## Generalizing the Matrix Multiplication Loop Nest

Matrix multiplication:

$$
\begin{aligned}
& \text { for } i=1: N \\
& \text { for } j=1: N \\
& \quad \text { for } k=1: N \\
& \quad \mathbf{C}_{i, j} \leftarrow \mathbf{C}_{i, j}+\mathbf{A}_{i, k} * \mathbf{B}_{k, j}
\end{aligned}
$$

| Parameter | Description | Example: Matrix Mult. |
| :---: | :---: | :---: |
| $d$ | depth of loop nest | 3 |

## Generalizing the Matrix Multiplication Loop Nest

Matrix multiplication:

$$
\begin{aligned}
& \text { for } i=1: N \\
& \text { for } j=1: N \\
& \quad \text { for } k=1: N \\
& \quad \mathbf{C}_{i, j} \leftarrow \mathbf{C}_{i, j}+\mathbf{A}_{i, k} * \mathbf{B}_{k, j}
\end{aligned}
$$

| Parameter | Description | Example: Matrix Mult. |
| :---: | :---: | :---: |
| $d$ | depth of loop nest | 3 |
| $Z$ | iteration space $\left(\right.$ subset of $\left.\mathbb{Z}^{d}\right)$ | $\{1, \ldots, N\}^{3}$ |

## Generalizing the Matrix Multiplication Loop Nest

Matrix multiplication:

$$
\begin{aligned}
& \text { for } i=1: N \\
& \text { for } j=1: N \\
& \quad \text { for } k=1: N \\
& \quad \mathbf{C}_{i, j} \leftarrow \mathbf{C}_{i, j}+\mathbf{A}_{i, k} * \mathbf{B}_{k, j}
\end{aligned}
$$

| Parameter | Description | Example: Matrix Mult. |
| :---: | :---: | :---: |
| $d$ | depth of loop nest | 3 |
| $Z$ | iteration space (subset of $\left.\mathbb{Z}^{d}\right)$ | $\{1, \ldots, N\}^{3}$ |
| $m$ | number of arrays | 3 |

## Generalizing the Matrix Multiplication Loop Nest

Matrix multiplication:

$$
\begin{aligned}
& \text { for } i=1: N \\
& \text { for } j=1: N \\
& \quad \text { for } k=1: N \\
& \quad \mathbf{C}_{i, j} \leftarrow \mathbf{C}_{i, j}+\mathbf{A}_{i, k} * \mathbf{B}_{k, j}
\end{aligned}
$$

| Parameter | Description | Example: Matrix Mult. |
| :---: | :---: | :---: |
| $d$ | depth of loop nest | 3 |
| $Z$ | iteration space (subset of $\left.\mathbb{Z}^{d}\right)$ | $\{1, \ldots, N\}^{3}$ |
| $m$ | number of arrays | 3 |
| $A_{1}, \ldots, A_{m}$ | arrays | A, B, C |

## Generalizing the Matrix Multiplication Loop Nest

Matrix multiplication:

$$
\begin{aligned}
& \text { for } i=1: N \\
& \text { for } j=1: N \\
& \quad \text { for } k=1: N \\
& \quad \mathbf{C}_{i, j} \leftarrow \mathbf{C}_{i, j}+\mathbf{A}_{i, k} * \mathbf{B}_{k, j}
\end{aligned}
$$

| Parameter | Description | Example: Matrix Mult. |
| :---: | :---: | :---: |
| $d$ | depth of loop nest | 3 |
| $Z$ | iteration space (subset of $\left.\mathbb{Z}^{d}\right)$ | $\{1, \ldots, N\}^{3}$ |
| $m$ | number of arrays | 3 |
| $A_{1}, \ldots, A_{m}$ | arrays | $\mathbf{A}, \mathbf{B}, \mathbf{C}$ |
| $d_{1}, \ldots, d_{m}$ | array dimensions | $2,2,2$ |

## Generalizing the Matrix Multiplication Loop Nest

Matrix multiplication:

$$
\begin{aligned}
& \text { for } i=1: N \\
& \text { for } j=1: N \\
& \quad \text { for } k=1: N \\
& \quad \mathbf{C}_{i, j} \leftarrow \mathbf{C}_{i, j}+\mathbf{A}_{i, k} * \mathbf{B}_{k, j}
\end{aligned}
$$

| Parameter | Description | Example: Matrix Mult. |
| :---: | :---: | :---: |
| $d$ | depth of loop nest | 3 |
| $Z$ | iteration space (subset of $\mathbb{Z}^{d}$ ) | $\{1, \ldots, N\}^{3}$ |
| $m$ | number of arrays | 3 |
| $A_{1}, \ldots, A_{m}$ | arrays | $\mathbf{A}, \mathbf{B}, \mathbf{C}$ |
| $d_{1}, \ldots, d_{m}$ | array dimensions | $2,2,2$ |
| $\phi_{1}, \ldots, \phi_{m}$ | array subscripts (linear maps: $\left.\mathbb{Z}^{d} \rightarrow \mathbb{Z}^{d_{j}}\right)$ | $(i, j, k) \mapsto(i, k),(k, j),(i, j)$ |

## Generalizing the Matrix Multiplication Loop Nest

Matrix multiplication:

$$
\begin{aligned}
& \text { for } i=1: N \\
& \text { for } j=1: N \\
& \quad \text { for } k=1: N \\
& \quad \mathbf{C}_{i, j} \leftarrow \mathbf{C}_{i, j}+\mathbf{A}_{i, k} * \mathbf{B}_{k, j}
\end{aligned}
$$

| Parameter | Description | Example: Matrix Mult. |
| :---: | :---: | :---: |
| $d$ | depth of loop nest | 3 |
| $Z$ | iteration space (subset of $\mathbb{Z}^{d}$ ) | $\{1, \ldots, N\}^{3}$ |
| $m$ | number of arrays | 3 |
| $A_{1}, \ldots, A_{m}$ | arrays | $\mathbf{A}, \mathbf{B}, \mathbf{C}$ |
| $d_{1}, \ldots, d_{m}$ | array dimensions | $2,2,2$ |
| $\phi_{1}, \ldots, \phi_{m}$ | array subscripts (linear maps: $\left.\mathbb{Z}^{d} \rightarrow \mathbb{Z}^{d_{j}}\right)$ | $(i, j, k) \mapsto(i, k),(k, j),(i, j)$ |

## General loop nest:

for $i=\left(i_{1}, \ldots, i_{d}\right) \in Z$ inner_loop $\left(A_{1}\left(\phi_{1}(i)\right), \ldots, A_{m}\left(\phi_{m}(i)\right)\right)$

## General Loop Nests: Main Result

## General loop nest:

$$
\begin{aligned}
& \text { for } i=\left(i_{1}, \ldots, i_{d}\right) \in Z \\
& \quad \text { inner_loop }\left(A_{1}\left(\phi_{1}(i)\right), \ldots, A_{m}\left(\phi_{m}(i)\right)\right)
\end{aligned}
$$

## Main Result: Communication Lower Bound for General Loop Nests

LB: There exists $\alpha \geq 0$ such that, in any implementation, $\mathcal{C} \succeq \frac{|Z|}{M^{\alpha}}$.
UB: There exists a tiled implementation such that $\mathcal{C} \preceq \frac{|Z|}{M^{\alpha}}$.

## General Loop Nests: Main Result

## General loop nest:

$$
\begin{aligned}
& \text { for } i=\left(i_{1}, \ldots, i_{d}\right) \in Z \\
& \quad \text { inner_loop }\left(A_{1}\left(\phi_{1}(i)\right), \ldots, A_{m}\left(\phi_{m}(i)\right)\right)
\end{aligned}
$$

## Main Result: Communication Lower Bound for General Loop Nests

LB: There exists $\alpha \geq 0$ such that, in any implementation, $\mathcal{C} \succeq \frac{|Z|}{M^{\alpha}}$.
UB: There exists a tiled implementation such that $\mathcal{C} \preceq \frac{|Z|}{M^{\alpha}}$.
Exponent $\alpha$ and tile params. are computable from the subscripts $\left(\phi_{j}\right)_{j}$.

- Exponent $\alpha$ is a rational number between 0 and $m-1$.
- Tiles are parallelotopes (in $\mathbb{Z}^{d}$ ) of size $\approx M^{\alpha+1}$.
- Matrix multiply: $\alpha=1 / 2$; tiles are cubes of size $\approx M^{3 / 2}$.


## New Ingredient: Hölder-Brascamp-Lieb (HBL) Inequality

Continuing notations $d, m, \phi_{1}, \ldots, \phi_{m}$ from the general loop nest,

## Hölder-Brascamp-Lieb (HBL) Inequality

If there exists $s=\left(s_{1}, \ldots, s_{m}\right) \in[0,1]^{m}$ such that,
for all subgroups $H$ of $\mathbb{Z}^{d}$,

$$
\operatorname{rank} H \leq \sum_{j=1}^{m} s_{j} \cdot \operatorname{rank} \phi_{j}(H),
$$

then,
for all subsets $S$ of $\mathbb{Z}^{d}$,

$$
|S| \leq \prod_{j=1}^{m}\left|\phi_{j}(S)\right|^{s_{j}}
$$

## New Ingredient: Hölder-Brascamp-Lieb (HBL) Inequality

Continuing notations $d, m, \phi_{1}, \ldots, \phi_{m}$ from the general loop nest,

## Hölder-Brascamp-Lieb (HBL) Inequality

If there exists $s=\left(s_{1}, \ldots, s_{m}\right) \in[0,1]^{m}$ such that,
for all subgroups $H$ of $\mathbb{Z}^{d}$,

$$
\operatorname{rank} H \leq \sum_{j=1}^{m} s_{j} \cdot \operatorname{rank} \phi_{j}(H),
$$

then,
for all subsets $S$ of $\mathbb{Z}^{d}$,

$$
|S| \leq \prod_{j=1}^{m}\left|\phi_{j}(S)\right|^{s_{j}}
$$

Loomis-Whitney inequality is a special case:

- $m=d=3$ and $\left(\phi_{1}, \phi_{2}, \phi_{3}\right)=\left(\phi_{\mathbf{A}}, \phi_{\mathbf{B}}, \phi_{\mathbf{C}}\right)$
- Hypothesis satisfied by $\left(s_{1}, s_{2}, s_{3}\right)=(1 / 2,1 / 2,1 / 2)$.


## HBL-LP: A Linear Program for the Optimal Exponent

- HBL lets us bound below the memory footprint of iterations $S$,

$$
\mu(S)=\sum_{j=1}^{m}\left|\phi_{j}(S)\right| \geq|S|^{1 / \sigma(s)}, \quad \text { abbreviating } \sigma(s)=\sum_{j=1}^{m} s_{j}
$$

## HBL-LP: A Linear Program for the Optimal Exponent

- HBL lets us bound below the memory footprint of iterations $S$,

$$
\mu(S)=\sum_{j=1}^{m}\left|\phi_{j}(S)\right| \geq|S|^{1 / \sigma(s)}, \quad \text { abbreviating } \sigma(s)=\sum_{j=1}^{m} s_{j}
$$

- The segmentation argument yields a communication lower bound,

$$
\mathcal{C} \succeq \frac{|Z|}{M^{\sigma(s)-1}}, \quad \text { meaning that the exponent } \alpha \leq \sigma(s)-1
$$

## HBL-LP: A Linear Program for the Optimal Exponent

- HBL lets us bound below the memory footprint of iterations $S$,

$$
\mu(S)=\sum_{j=1}^{m}\left|\phi_{j}(S)\right| \geq|S|^{1 / \sigma(s)}, \quad \text { abbreviating } \sigma(s)=\sum_{j=1}^{m} s_{j}
$$

- The segmentation argument yields a communication lower bound,

$$
\mathcal{C} \succeq \frac{|Z|}{M^{\sigma(s)-1}}, \quad \text { meaning that the exponent } \alpha \leq \sigma(s)-1
$$

- A lower bound of this form exists for any $s \in \mathcal{P}$, where

$$
\mathcal{P}=\left\{s \in[0,1]^{m} \mid \quad\left(\forall H \leq \mathbb{Z}^{d}\right) \text { rank } H \leq \sum_{j=1}^{m} s_{j} \cdot \operatorname{rank} \phi_{j}(H)\right\} ;
$$

to find the sharpest such bound, we with to solve a linear program,

$$
\alpha+1=\min \{\sigma(s) \mid s \in \mathcal{P}\}
$$

## HBL-LP: Challenges and a Solution Algorithm

(Bad news) Checking membership in $\mathcal{P}$ nominally requires considering all (infinitely many) subgroups $H \leq \mathbb{Z}^{d}$.
$\mathcal{P}$ for matrix multiplication:


## HBL-LP: Challenges and a Solution Algorithm

(Bad news) Checking membership in $\mathcal{P}$ nominally requires considering all (infinitely many) subgroups $H \leq \mathbb{Z}^{d}$.
(Good news) The subgroup ranks are integers between 0 and $d$ : the linear constraints defining $\mathcal{P}$ are elements of a finite set.
$\mathcal{P}$ for matrix multiplication:


## HBL-LP: Challenges and a Solution Algorithm

(Bad news) Checking membership in $\mathcal{P}$ nominally requires considering all (infinitely many) subgroups $H \leq \mathbb{Z}^{d}$.
(Good news) The subgroup ranks are integers between 0 and $d$ : the linear constraints defining $\mathcal{P}$ are elements of a finite set.
(Bad news?) The existence of an algorithm that decides whether a particular constraint is involved in the HBL-LP is equivalent to Hilbert's Tenth Problem over $\mathbb{Q}$, conjectured to have a negative answer.
$\mathcal{P}$ for matrix multiplication:


## HBL-LP: Challenges and a Solution Algorithm

(Bad news) Checking membership in $\mathcal{P}$ nominally requires considering all (infinitely many) subgroups $H \leq \mathbb{Z}^{d}$.
(Good news) The subgroup ranks are integers between 0 and $d$ : the linear constraints defining $\mathcal{P}$ are elements of a finite set.
(Bad news?) The existence of an algorithm that decides whether a particular constraint is involved in the HBL-LP is equivalent to Hilbert's Tenth Problem over $\mathbb{Q}$, conjectured to have a negative answer.
$\mathcal{P}$ for matrix multiplication:


## Good News

Membership in $\mathcal{P}$ is decidable.

## Tiling with Parallelotopes

## Existence of an Optimal Tiling

There exist parallelotopes $T$ satisfying $|T| \succeq M^{\alpha+1}$ and $\mu(T) \leq M$.
Ingredients: linear duality, subgroup series (flags), Smith normal form.

## Tiling with Parallelotopes

## Existence of an Optimal Tiling

There exist parallelotopes $T$ satisfying $|T| \succeq M^{\alpha+1}$ and $\mu(T) \leq M$.
Ingredients: linear duality, subgroup series (flags), Smith normal form.


## Tiling with Parallelotopes

## Existence of an Optimal Tiling

There exist parallelotopes $T$ satisfying $|T| \succeq M^{\alpha+1}$ and $\mu(T) \leq M$.
Ingredients: linear duality, subgroup series (flags), Smith normal form.


## Example: Matrix Multiplication (Revisited)

Loop nest:
for $i=1: N$

$$
\begin{aligned}
& \text { for } j=1: N \\
& \text { for } k=1: N \\
& \quad \mathbf{C}_{i, j} \leftarrow \mathbf{C}_{i, j}+\mathbf{A}_{i, k} * \mathbf{B}_{k, j}
\end{aligned}
$$

Parameters:
$d=3, \quad Z=\{1, \ldots, N\}^{3}, \quad m=3$,
$(i, j, k) \xrightarrow{\phi_{1}}(i, k) \xrightarrow{\stackrel{A_{1}}{\longmapsto}} \mathbf{A}_{i, k}$
$(i, j, k) \xrightarrow{\stackrel{\phi_{2}}{\longmapsto}}(k, j) \xrightarrow{A_{2}} \mathbf{B}_{k, j}$
$(i, j, k) \xrightarrow{\phi_{3}}(i, j) \xrightarrow{A_{3}} \mathbf{C}_{i, j}$

## Example: Matrix Multiplication (Revisited)

Loop nest:

$$
\begin{aligned}
& \text { for } i=1: N \\
& \text { for } j=1: N \\
& \quad \text { for } k=1: N \\
& \quad \mathbf{C}_{i, j} \leftarrow \mathbf{C}_{i, j}+\mathbf{A}_{i, k} * \mathbf{B}_{k, j}
\end{aligned}
$$

Parameters:

$$
\begin{aligned}
& d=3, \quad Z=\{1, \ldots, N\}^{3}, \quad m=3, \\
& (i, j, k) \xrightarrow{\stackrel{\phi_{1}}{\longmapsto}}(i, k) \xrightarrow{\stackrel{A_{1}}{\longmapsto}} \mathbf{A}_{i, k} \\
& (i, j, k) \xrightarrow{\phi_{2}}(k, j) \xrightarrow{A_{2}} \mathbf{B}_{k, j} \\
& (i, j, k) \xrightarrow{\phi_{3}}(i, j) \xrightarrow{A_{3}} \mathbf{C}_{i, j}
\end{aligned}
$$

HBL theory gives:

- Exponent $\alpha=1 / 2$ (data reuse $\preceq M^{1 / 2}$ );
- Tile $T=\{1, \ldots, b\}^{3}$, with $b \approx M^{1 / 2}$.


## Communication Lower Bound

The lower bound $\mathcal{C} \succeq \frac{N^{3}}{M^{1 / 2}}$ is attainable by tiling with translates of $T$.

## Example: Matrix-Vector Multiplication

Loop nest:

$$
\begin{aligned}
& \text { for } i=1: N \\
& \quad \text { for } j=1: N \\
& \quad \mathbf{y}_{i} \leftarrow \mathbf{y}_{i}+\mathbf{A}_{i, j} * \mathbf{x}_{j}
\end{aligned}
$$

Parameters:
$d=2, \quad Z=\{1, \ldots, N\}^{2}, \quad m=3$,

$$
\begin{array}{llcl}
(i, j) & \xrightarrow[\phi_{1}]{\longmapsto} & (i, j) & \xrightarrow[A_{1}]{\longmapsto} \\
(i, j) & \mathbf{A}_{i, j} \\
(i, j) & \xrightarrow[\phi_{2}]{\longmapsto} & (j) & \xrightarrow{A_{3}} \\
\longmapsto & (i) & \xrightarrow{A_{3}} & \mathbf{x}_{j} \\
\longmapsto & \mathbf{y}_{i}
\end{array}
$$

## Example: Matrix-Vector Multiplication

Loop nest:

$$
\begin{aligned}
& \text { for } i=1: N \\
& \quad \text { for } j=1: N \\
& \quad \mathbf{y}_{i} \leftarrow \mathbf{y}_{i}+\mathbf{A}_{i, j} * \mathbf{x}_{j}
\end{aligned}
$$

Parameters:

$$
\begin{aligned}
& d=2, \quad Z=\{1, \ldots, N\}^{2}, \quad m=3, \\
& (i, j) \xrightarrow{\phi_{1}}(i, j) \xrightarrow{A_{1}} \mathbf{A}_{i, j} \\
& (i, j) \xrightarrow{\phi_{2}}(j) \xrightarrow{A_{2}} \mathbf{x}_{j} \\
& (i, j) \xrightarrow{\phi_{3}}(i) \xrightarrow{A_{3}} \quad \mathbf{y}_{i}
\end{aligned}
$$

HBL theory gives:

- Exponent $\alpha=0$ (data reuse $\preceq 1$ );
- Tile $T=\{(1,1,1)\}^{3}$.


## Communication Lower Bound

The lower bound $\mathcal{C} \succeq N^{2}$ is attainable by tiling with translates of $T$.

## Example: Tensor Contraction

## Loop nest:

$$
\begin{aligned}
& \text { for } i=1: N \\
& \text { for } j=1: N \\
& \text { for } k=1: N \\
& \text { for } I=1: N \\
& \quad \text { for } m=1: N \\
& \quad \mathbf{C}_{i, j, k} \leftarrow \mathbf{C}_{i, j, k}+\mathbf{A}_{i, l, m} * \mathbf{B}_{l, m, j, k}
\end{aligned}
$$

## Parameters:

$$
\begin{aligned}
& d=5, \quad Z=\{1, \ldots, N\}^{5}, \quad m=3, \\
& (i, j, k, l, m) \xrightarrow{\phi_{1}} \quad(i, l, m) \quad \stackrel{A_{1}}{\longmapsto} \quad \mathbf{A}_{i, l, m} \\
& (i, j, k, l, m) \xrightarrow{\phi_{2}}(I, m, j, k) \xrightarrow{A_{2}} \mathbf{B}_{l, m, j, k} \\
& (i, j, k, l, m) \xrightarrow{\phi_{3}} \quad(i, j, k) \quad \stackrel{A_{3}}{\longmapsto} \mathbf{C}_{i, j, k}
\end{aligned}
$$

## Example: Tensor Contraction

## Loop nest:

$$
\begin{aligned}
& \text { for } i=1: N \\
& \text { for } j=1: N \\
& \text { for } k=1: N \\
& \quad \text { for } I=1: N \\
& \quad \text { for } m=1: N \\
& \quad \mathbf{C}_{i, j, k} \leftarrow \mathbf{C}_{i, j, k}+\mathbf{A}_{i, l, m} * \mathbf{B}_{l, m, j, k}
\end{aligned}
$$

Parameters:

$$
\begin{aligned}
& d=5, \quad Z=\{1, \ldots, N\}^{5}, \quad m=3, \\
& (i, j, k, l, m) \xrightarrow{\phi_{1}} \quad(i, l, m) \quad \stackrel{A_{1}}{\longmapsto} \quad \mathbf{A}_{i, l, m} \\
& (i, j, k, l, m) \xrightarrow{\phi_{2}}(l, m, j, k) \xrightarrow{A_{2}} \mathbf{B}_{l, m, j, k} \\
& (i, j, k, l, m) \xrightarrow{\phi_{3}} \quad(i, j, k) \quad \stackrel{A_{3}}{\longmapsto} \mathbf{C}_{i, j, k}
\end{aligned}
$$

HBL theory gives:

- Exponent $\alpha=1 / 2$ (data reuse $\preceq M^{1 / 2}$, same as matrix multiply!);
- Tile $T=X_{i=1}^{5}\left\{1, \ldots, b_{i}\right\}$ with $b_{1}, b_{2} b_{3}, b_{4} b_{5} \approx M^{1 / 2}$.


## Communication Lower Bound

The lower bound $\mathcal{C} \succeq \frac{N^{5}}{M^{1 / 2}}$ is attainable by tiling with translates of $T$.

## Example: Particle Simulation

Loop nest:

$$
\begin{aligned}
& \text { for } i=1: N \\
& \quad \text { for } j=1: N \\
& \quad \mathbf{V}_{i} \leftarrow \mathbf{V}_{i}+G\left(\mathbf{P}_{i}, \mathbf{P}_{j}\right)
\end{aligned}
$$

## Parameters:

$$
d=2, \quad Z=\{1, \ldots, N\}^{2}, \quad m=3
$$

$$
(i, j) \stackrel{\phi_{1}}{\longleftrightarrow}(i) \xrightarrow{\stackrel{A_{1}}{\longleftrightarrow}} \mathbf{P}_{i}
$$

$$
(i, j) \xrightarrow{\phi_{2}}(j) \xrightarrow{\stackrel{A_{2}}{\longrightarrow}} \mathbf{P}_{j}
$$

$$
(i, j) \xrightarrow{\phi_{3}}(i) \xrightarrow{A_{3}} \mathbf{V}_{i}
$$

## Example: Particle Simulation

Loop nest:

$$
\begin{aligned}
& \text { for } i=1: N \\
& \quad \text { for } j=1: N \\
& \quad \mathbf{V}_{i} \leftarrow \mathbf{V}_{i}+G\left(\mathbf{P}_{i}, \mathbf{P}_{j}\right)
\end{aligned}
$$

## Parameters:

$$
\begin{aligned}
& d=2, \quad Z=\{1, \ldots, N\}^{2}, \quad m=3, \\
& (i, j) \xrightarrow{\phi_{1}}(i) \xrightarrow{\longmapsto} \xrightarrow{A_{1}} \mathbf{P}_{i} \\
& (i, j) \xrightarrow{\phi_{2}}(j) \xrightarrow{A_{2}} \mathbf{P}_{j} \\
& (i, j) \xrightarrow{\phi_{3}}(i) \xrightarrow{A_{3}} \mathbf{V}_{i}
\end{aligned}
$$

HBL theory gives:

- Exponent $\alpha=1$ (data reuse $\preceq M$ );
- Tile $T=\{1, \ldots, b\}^{2}$, with $b \approx M$.


## Communication Lower Bound

The lower bound $\mathcal{C} \succeq \frac{N^{2}}{M}$ is attainable by tiling with translates of $T$.

## Example: Convolutional Neural Network

Loop nest:

```
for \(i=1: N\)
    for \(j=1: N\)
        for \(k=1: N\)
            for \(I=1: N\)
            for \(m=1: N\)
                for \(n=1: N\)
                    \(\mathbf{C}_{i, j, m, n} \leftarrow \mathbf{C}_{i, j, m, n}+\mathbf{A}_{i, k, l} * \mathbf{B}_{j, k+m, l+n}\)
```


## Parameters:

$$
\begin{aligned}
& d=6, \quad Z=\{1, \ldots, N\}^{6}, \quad m=3, \\
& (i, j, k, l, m, n) \xrightarrow{\phi_{1}} \quad(i, k, l) \quad \xrightarrow{A_{1}} \quad \mathbf{A}_{i, k, l} \\
& (i, j, k, l, m, n) \xrightarrow{\phi_{2}}(j, k+m, l+n) \xrightarrow{A_{2}} \mathbf{B}_{j, k+m, l+n} \\
& (i, j, k, l, m, n) \xrightarrow{\phi_{3}} \quad(i, j, m, n) \quad \stackrel{A_{3}}{\longmapsto} \quad \mathbf{C}_{i, j, m, n}
\end{aligned}
$$

## Example: Convolutional Neural Network

Loop nest:

$$
\begin{aligned}
& \text { for } i=1: N \\
& \text { for } j=1: N \\
& \text { for } k=1: N \\
& \text { for } I=1: N \\
& \quad \text { for } m=1: N \\
& \quad \text { for } n=1: N \\
& \quad \mathbf{C}_{i, j, m, n} \leftarrow \mathbf{C}_{i, j, m, n}+\mathbf{A}_{i, k, l} * \mathbf{B}_{j, k+m, l+n}
\end{aligned}
$$

Parameters:

$$
\begin{aligned}
& d=6, \quad Z=\{1, \ldots, N\}^{6}, \quad m=3, \\
& (i, j, k, l, m, n) \xrightarrow{\phi_{1}} \quad(i, k, l) \quad \xrightarrow{\stackrel{A_{1}}{\longrightarrow}} \quad \mathbf{A}_{i, k, l} \\
& (i, j, k, l, m, n) \xrightarrow{\phi_{2}}(j, k+m, l+n) \xrightarrow{A_{2}} \mathbf{B}_{j, k+m, l+n} \\
& (i, j, k, l, m, n) \xrightarrow{\phi_{3}} \quad(i, j, m, n) \quad \stackrel{A_{3}}{\longmapsto} \quad \mathbf{C}_{i, j, m, n}
\end{aligned}
$$

HBL theory gives:

- Exponent $\alpha=1$ (data reuse $\preceq M$, same as particle simulation);
- Tile $T=X_{i=1}^{6}\left\{1, \ldots, b_{i}\right\}$ with $b_{1}, b_{2} \approx 1, b_{3}, b_{4}, b_{5}, b_{6} \approx M^{1 / 2}$.


## Communication Lower Bound

The lower bound $\mathcal{C} \succeq \frac{N^{6}}{M}$ is attainable by tiling with translates of $T$.

## Concluding Remarks

- Communication lower bounds for loop nests


## Concluding Remarks

- Communication lower bounds for loop nests
- Lower bounds attainable by iteration space tiling


## Concluding Remarks

- Communication lower bounds for loop nests
- Lower bounds attainable by iteration space tiling
- Goal: compiler automatically generates optimal code


## Concluding Remarks

- Communication lower bounds for loop nests
- Lower bounds attainable by iteration space tiling
- Goal: compiler automatically generates optimal code

Many details omitted:

- Parallel implementations: replace $|Z|$ by $|Z| / P$


## Concluding Remarks

- Communication lower bounds for loop nests
- Lower bounds attainable by iteration space tiling
- Goal: compiler automatically generates optimal code

Many details omitted:

- Parallel implementations: replace $|Z|$ by $|Z| / P$
- Constants hidden in asymptotic notation $(\preceq, \approx, \succeq)$


## Concluding Remarks

- Communication lower bounds for loop nests
- Lower bounds attainable by iteration space tiling
- Goal: compiler automatically generates optimal code

Many details omitted:

- Parallel implementations: replace $|Z|$ by $|Z| / P$
- Constants hidden in asymptotic notation $(\preceq, \approx, \succeq)$
- Attainability when $Z$ is irregular (e.g., sparse)


## Concluding Remarks

- Communication lower bounds for loop nests
- Lower bounds attainable by iteration space tiling
- Goal: compiler automatically generates optimal code

Many details omitted:

- Parallel implementations: replace $|Z|$ by $|Z| / P$
- Constants hidden in asymptotic notation $(\preceq, \approx, \succeq)$
- Attainability when $Z$ is irregular (e.g., sparse)
- Attainability when data dependences


## Concluding Remarks

- Communication lower bounds for loop nests
- Lower bounds attainable by iteration space tiling
- Goal: compiler automatically generates optimal code

Many details omitted:

- Parallel implementations: replace $|Z|$ by $|Z| / P$
- Constants hidden in asymptotic notation $(\preceq, \approx, \succeq)$
- Attainability when $Z$ is irregular (e.g., sparse)
- Attainability when data dependences
- Complexity of computing lower bound and optimal tiling


## Concluding Remarks

- Communication lower bounds for loop nests
- Lower bounds attainable by iteration space tiling
- Goal: compiler automatically generates optimal code

Many details omitted:

- Parallel implementations: replace $|Z|$ by $|Z| / P$
- Constants hidden in asymptotic notation $(\preceq, \approx, \succeq)$
- Attainability when $Z$ is irregular (e.g., sparse)
- Attainability when data dependences
- Complexity of computing lower bound and optimal tiling
- Optimizing subprograms, compositions of programs


## Concluding Remarks

- Communication lower bounds for loop nests
- Lower bounds attainable by iteration space tiling
- Goal: compiler automatically generates optimal code

Many details omitted:

- Parallel implementations: replace $|Z|$ by $|Z| / P$
- Constants hidden in asymptotic notation $(\preceq, \approx, \succeq)$
- Attainability when $Z$ is irregular (e.g., sparse)
- Attainability when data dependences
- Complexity of computing lower bound and optimal tiling
- Optimizing subprograms, compositions of programs

Please see tech. reports UCB/EECS-2013-61 and UCB/EECS-2015-185.

## Thank You!

