Local Fourier Uniformity

July 24, 2019

The main question

Question: Is the multiplicative and additive structure of the integers independent?
Conjectural answer: Yes, up to minor local obstructions.
The precise form of this answer is due to Chowla and Elliott.

Conjecture (Chowla-Elliott)

Let $f_1, f_2, \ldots, f_r$ be real-valued multiplicative functions with $|f_i| \leq 1$ (i.e $f_i(mn) = f_i(m)f_i(n)$ for m, n co-prime). Then, for any distinct $h_1, \ldots, h_r$ there exists a constant $C_{h_1,\ldots,h_r}$ such that ,

$$\frac{1}{x}\sum_{n\leq x} f_1(n+h_1)\ldots f_r(n+h_r) \sim C_{h_1,\ldots,h_r} \prod_{i=1}^{r} \Big(\frac{1}{x}\sum_{n\leq x} f_i(n)\Big).$$

Note that $C_{h_1,\ldots,h_r}$ can be zero (this is in fact the most interesting case!) and in that case we interpret the asymptotic as saying that the left-hand side is $o(1)$.

Special case

1. A special case of particular interest is $f_1 = \ldots = f_r = \lambda$ with $\lambda(n) = (-1)^{\Omega(n)}$ the Liouville function, where $\Omega(n)$ is the number of prime factors of n.
2. The reason for this is that the case $r = 1$, $f_1 = \lambda$, i.e

   $$\sum_{n \leq x} \lambda(n) = o(x)$$

   is equivalent to the prime number theorem.
3. So the problem of establishing, for distinct $h_1, \ldots, h_r$,

   $$\sum_{n \leq x} \lambda(n)\lambda(n+h_1)\ldots\lambda(n+h_r) = o(x)$$

   is viewed as analogous to a prime number theorem for twin primes or more generally tuplets of primes (but it does not directly imply it).

Technical motivation

1. A lot of the technical motivation comes from the shifted convolution problem (SCP),
$$\sum_{n \leq x} \lambda_\pi(n)\lambda_\pi(n+h)$$
which is important in the analytic theory of L-functions.
2. The shifted convolution problem is completely solved for π automorphic forms of rank 1 and 2.
3. It is completely open for higher rank.
4. The main issue in the SCP is the same as in the Chowla-Elliott conjecture : understanding the multiplicative correlations of nearby integers.
5. However the SCP is often addressed by completely ignoring multiplicativity and only using automorphy (in my opinion this is part of the reason we are stuck in higher rank).

Disclaimer

1. Until recently the Chowla-Elliott conjecture also appeared to be completely out of reach
2. Since 2015 there has been a lot of progress.
3. I am (overly?) optimistic that we will see a resolution within the next 10 years and that this will spill over into progress on the SCP.
4. Since there has been rapid progress there are some patchy "gray areas" which in principle ought to be not there, and we don't understand why they are there. I will therefore try to give a quick survey of the state of the art (nonetheless I will omit some results).

Recent progress (Averaged versions)

There has been substantial recent progress on "averaged" versions of the Chowla-Elliott conjecture.

Theorem (Matomäki-Radziwiłł, 2015)

Let $f_1, f_2 : \mathbb{N} \to \mathbb{R}$ be multiplicative functions with $|f_1|, |f_2| \leq 1$. Then, for any $H \to \infty$ arbitrarily slowly with $x \to \infty$,

$$\frac{1}{2H} \sum_{|h| \leq H} \frac{1}{x} \sum_{n \leq x} f_1(n) f_2(n+h) \sim C \prod_{i=1}^{2} \Big(\frac{1}{x} \sum_{n \leq x} f_i(n) \Big).$$

with $C > 0$ a constant.

1. Actually we obtained a description of the behavior of short averages,

$$\sum_{x \leq n \leq x+H} f(n)$$

   for almost all $x \in [X, 2X]$. This essentially implies the above statement.

2. We can also prove variants in the case when f_i are coefficients of arbitrary rank (Hecke) automorphic forms.

Recent progress (Averaged versions)

Theorem (Matomäki-Radziwiłł-Tao, 2015)

Let $f_1, f_2 : \mathbb{N} \to \mathbb{R}$ be multiplicative functions with $|f_1|, |f_2| \leq 1$. Then, for any $H \to \infty$ arbitrarily slowly with $x \to \infty$,

$$\frac{1}{2H} \sum_{|h| \leq H} \Big| \frac{1}{x} \sum_{n \leq x} f_1(n) f_2(n+h) - C_h \prod_{i=1}^{2} \Big(\frac{1}{x} \sum_{n \leq x} f_i(n) \Big) \Big| = o(1)$$

for some constants $C_h = O(1)$.

In the proof one needs control not only over,

$$\sum_{x \leq n \leq x+H} f(n) \text{ but also over } \sum_{x \leq n \leq x+H} f(n) e(n\alpha)$$

for almost all $x \in [X, 2X]$ and with α fixed. Roughly speaking the case of $\alpha \in \mathbb{Q}$ follows from my result with Matomäki, while the case of $\alpha \notin \mathbb{Q}$ from a version of Vinogradov's method due to Daboussi-Delange-Katai-Bourgain-Sarnak-Ziegler. **Note:** Can also prove variants for coefficients of arbitrary rank (Hecke) automorphic forms

Recent progress (Logarithmic versions)

Theorem (Tao, 2015)

Let $f_1, f_2 : \mathbb{N} \to \mathbb{R}$ *be multiplicative functions. Then, given* h*, there exists a* C_h *(possibly zero) such that as* $x \to \infty$*,*

$$\sum_{n \leq x} \frac{f_1(n) f_2(n+h)}{n} \sim C_h \log x$$

1. The proof depends on the previous two results and the "entropy decrement argument". The latter is very sensitive to f_1, f_2 being non-lacunary and mostly of size ≈ 1 (in absolute value).
2. Because of this there are no analogues for f_i coefficients of automorphic forms.
3. The logarithmic averaging allows one to introduce a third variable, which allows to make use of the previous "averaged" results.

Recent progress (Logarithmic versions)

Let $\lambda(n)$ denote the Liouville function (i.e $\lambda(n) = (-1)^{\Omega(n)}$ where $\Omega(n)$ is the number of prime factors of n counted with multiplicity).

Theorem (Tao-Teräväinen, 2017)

Let k be **odd**. *Then, for any $h_1, \dots, h_k$ distinct, as $x \to \infty$,*

$$\sum_{n \le x} \frac{\lambda(n+h_1)\dots\lambda(n+h_k)}{n} = o(\log x)$$

1. In some sense this is a "parity trick". The proof depends on the fact that k is odd and $\lambda(p) = -1$.
2. The number theoretic content is much lighter (in comparison with the $k = 2$ case). For example for $k = 3$ one only needs cancellations in

$$\sum_{n \le x} \lambda(n) e(n\alpha).$$

which was proven by Davenport using ideas of Vinogradov.

Major questions

1. Can we show that for **any** $k \geq 2$, and any distinct $h_1, \ldots, h_k$,

$$\sum_{n \leq x} \frac{\lambda(n+h_1) \ldots \lambda(n+h_k)}{n} = o(\log x) \text{ ?} \tag{1}$$

   This would, among other things, settle Sarnak's conjecture in logarithmic form

2. Can we establish (1) without the logarithmic weights? (Even only for $k = 2$?)
3. Can we establish (1) for coefficients of automorphic forms?

The roadmap towards 1. is at the moment much more clear than the one towards 2. (There is some progress on 2. by Tao-Teräväinen, but in my opinion there is still no clear roadmap for 2.) Progress on 3. would require an alternative to the entropy decrement argument and at present there are no viable alternatives.

Fourier uniformity

Tao showed that in order to obtain

$$\sum_{n \le x} \frac{\lambda(n+h_1)\dots\lambda(n+h_k)}{n} = o(\log x) \qquad (2)$$

for every $k \geq 1$ it is necessary to establish the following

Conjecture (Local Fourier Uniformity conjecture)

Let $k \geq 1$ be given. For $H \to \infty$ arbitrarily slowly with $X \to \infty$,

$$\int_1^X \sup_{\substack{P \in \mathbb{R}[X] \\ deg\ P = k}} \Big| \sum_{x \le n \le x+H} \lambda(n) e(P(n)) \Big| \cdot \frac{dx}{x} = o(\log X)$$

1. In reality one also needs to establish the above for nilsequences. This is then also sufficient for (2).
2. As far as I can see the measure $\frac{dx}{x}$ appears to give no real advantage compared to dx. We will therefore consider the measure dx instead.

Previous results

Previous results are rather unsatisfactory:

Theorem (The $k = 0$ case, MR, 2015)

Let $H \to \infty$ arbitrarily slowly with $X \to \infty$. Then,

$$\int_1^X \Big| \sum_{x \leq n \leq x+H} \lambda(n) \Big| dx = o(HX).$$

This is a restatement of my earlier results with Matomäki.

Theorem (The $k = 1$ case for $H > X^{5/8}$, Zhan, 1991)

Let $\varepsilon > 0$. Let $H > X^{5/8+\varepsilon}$. Then, for $X \to \infty$,

$$\sum_{X \leq n \leq X+H} \lambda(n) e(n\alpha) = o(H).$$

The above uses Heath-Brown's identity. The method hits a hard limit at $H = \sqrt{X}$ because it is based on Dirichlet polynomial techniques.

New results

Theorem (The $k = 1$ case for $H > X^{\varepsilon}$, MRT, 2018)

Let $\varepsilon > 0$ be given. Then for $H > X^{\varepsilon}$ as $X \to \infty$,

$$\int_1^X \sup_{\alpha} \Big| \sum_{x \le n \le x+H} \lambda(n) e(n\alpha) \Big| dx = o(HX).$$

1. One can prove a variant for general multiplicative functions (even unbounded). The function shouldn't be close to $\chi(n) n^{iT}$ with χ of bounded conductor and $T \ll X^2$.
2. It appears to be possible to lower H to $H > \exp((\log X)^{1/2+\varepsilon})$.

Work in progress

Progress at the American Institute for Mathematics :

1. After conversations with Teräväinen we understood that our proof extends fairly easily to polynomials. Thus we can get, for any $k \in \mathbb{N}$,

$$\int_1^X \sup_{\substack{P \in \mathbb{R}[X] \\ \deg P = k}} \Big| \sum_{x \leq n \leq x+H} \lambda(n) e(P(n)) \Big| dx = o(HX)$$

   for $H > \exp((\log X)^{1/2+\varepsilon})$.

2. As part of an on-going large AIM collaboration it appears that we can also handle the case of nilsequences for $H > \exp((\log X)^{1/2+\varepsilon})$.
3. So the current bottle-neck towards a full resolution of Chowla-Elliott in logarithmic form is the reduction of H from $\exp(\sqrt{\log X})$ to H growing arbitrarily slowly.

Consequences of Fourier Uniformity

A good illustration of the content of Fourier uniformity for $k = 1$ is:

Corollary (MRT, 2018)

Let $\varepsilon > 0$. Let $\alpha(n)$ and $\beta(n)$ be two **arbitrary** *sequences of complex numbers with $|\alpha(i)| \leq 1$ and $|\beta(i)| \leq 1$ for every $i \geq 1$. Then, for $H > X^{\varepsilon}$ as $X \to \infty$,*

$$\sum_{|h| \leq H} \sum_{n \leq X} \lambda(n)\alpha(n+h)\beta(n+2h) = o(HX).$$

Proof sketch: The above is roughly, H^{-1} times

$$\int_X^{2X} \int_0^1 \Big(\sum_{x \leq n \leq x+H} \lambda(n)e(n\alpha) \Big) \Big(\sum_{x \leq n \leq x+H} \alpha(n)e(n\alpha) \Big) \Big(\sum_{x \leq n \leq x+H} \beta(n)e(-2n\alpha) \Big)$$

Taking the supremum over α in the sum over $\lambda(n)$ and using Cauchy-Schwarz and Plancherel on the remaining two trigonometric polynomials we can bound the inner integral in absolute value by

$\sup_{\alpha} \Big| \sum_{x \leq n \leq x+H} \lambda(n)e(n\alpha) \Big| \cdot H$ and the result follows from Fourier uniformity

Consequences of Fourier uniformity

If the sequences $\alpha(n)$, $\beta(n)$ admit tight sieve majorants then we can allow them to be unbounded.

Corollary (MRT, 2018)

Let $\varepsilon > 0$ and $H > X^{\varepsilon}$. Then, as $X \to \infty$,

$$\sum_{|h| \le H} \sum_{n \le X} \lambda(n)\Lambda(n+h)\Lambda(n+2h) = o(HX). \tag{3}$$

1. We are not able to establish that,

$$\sum_{|h| \le H} \sum_{n \le X} \Lambda(n+h)\Lambda(n+2h) \sim HX \tag{4}$$

2. (4) is equivalent to a prime number theorem in almost all intervals of length H.
3. The best known result for (4) remains $H > X^{1/6+\varepsilon}$ due to Huxley (Zaccagnini showed using ideas of Heath-Brown that ε can tend to zero from the negative side).
4. In (3) all the heavy-lifting is done by the Liouville function

Consequences of Fourier uniformity : Work in progress

It is possible to generalize the result on Fourier uniformity to unbounded functions. This then has consequences for triple correlations of divisor functions (or more general coefficients of high-rank automorphic forms)

Corollary (MRT, 2019+)

Let $\varepsilon > 0$*. Let* $k, \ell, m \geq 1$*. Let* $H > X^{\varepsilon}$*. Then, as* $X \to \infty$*,*

$$\sum_{|h| \leq H} \sum_{n \leq X} d_k(n) d_\ell(n+h) d_m(n+2h) \sim CHX(\log X)^{k+\ell+m-3}.$$

with $C > 0$ *a constant.*

1. For $H = X$ this follows from work of Mathiessen.
2. In the case $k = \ell = m = 2$ Blomer obtained an asymptotic for $H > X^{1/3+\varepsilon}$ using spectral methods of automorphic forms.
3. It is striking that one can go further by using only multiplicativity (and the Littlewood zero-free region).
4. In many applications of spectral methods multiplicativity is never used (this is especially clear when variants apply to half-integral weight forms).

Sketch of proof of Fourier Uniformity for $k = 1$

1. I will now sketch the ideas going into the proof of the $k = 1$ case, with $H > X^{\varepsilon}$,
$$\int_1^X \sup_{\alpha} \Big| \sum_{x \le n \le x+H} \lambda(n) e(n\alpha) \Big| dx = o(HX).$$
2. We will suppose that Fourier Uniformity fails and reach a contradiction.
3. Arguing by contradiction we assume that there exists an $\eta > 0$ and a collection $\mathcal{I}$ of $\gg X/H$ disjoint intervals of length H, contained in $[1, X]$, such that, for each $I \in \mathcal{I}$,
$$\Big| \sum_{n \in I} \lambda(n) e(n\alpha_I) \Big| \ge \eta H$$
for some $\alpha_I \in \mathbb{R}$.

Overall strategy

1. We will gain increasing control over the frequencies α_I.
2. We will first show, that the multiplicativity of λ forces that for a positive proportion of I, $e(n\alpha_I) \approx n^{iT_I}$ for some $|T_I| \ll X^2/H$ depending on I (in reality $e(n\alpha_I) \approx e(na/q)n^{iT_I}$ with $q \ll X^\varepsilon$ but for simplicity let's ignore the q-aspect).
3. Importantly the values T_I and T_J are sometimes related. If $\frac{I}{p} \cap \frac{J}{q} \neq \emptyset$ for some primes $p, q \in [P, 2P]$ then $T_I = T_J$. P is roughly of size H^ε. This is also a consequence of multiplicativity.
4. We then show that the above relationships between the T_I's imply that there is a global T such that $T_I = T$ for a positive proportion of the interval I's. This step uses cancellations in $\sum_{H \le p \le 2H} p^{it}$.
5. Therefore for a positive proportion of intervals I,

$$\Big| \sum_{n \in I} \lambda(n) e(n\alpha_I) \Big| \approx \Big| \sum_{n \in I} \lambda(n) n^{iT} \Big| \gg H$$

and this is ruled out by my result with Matomäki.

Overall strategy

1. The first two steps are "local" : they use the large sieve, the Turan-Kubilius inequality and multiplicativity at the primes $p \leq H$.
2. The second step is "global" : it uses cancellations in $\sum_{H \leq p \leq 2H} p^{it}$.

I will start with the second step.

The “global” step

1. We assume that there is a P, of size roughly H^ε and a set of $\gg X/H$ disjoint intervals I, J and real numbers T_I, T_J, such that whenever
$$\frac{I}{p} \cap \frac{J}{q} \neq \emptyset$$
for some $p, q \in [P, 2P]$ then we have $T_I = T_J$.
2. We want to show that there exists a T such that for a positive proportion of intervals I we have $T_I = T$.
3. We can think of the above as a graph with vertices corresponding to intervals I, J and edges between I, J when there exists primes $p, q \in [P, 2P]$ such that $\frac{I}{p} \cap \frac{J}{q} \neq \emptyset$.
4. We want to show that this is an expander graph, so that there is one giant connected component, and so there is a universal T with $T_I = T$ for a positive proportion of intervals I.

The “global” step

1. Avoiding all technical details, the existence of a large component follows from this graph having a spectral gap.
2. This in turn follows from cancellations in an exponential sum over primes, i.e
$$\sum_{H \le p \le 2H} p^{it} = o\left(\frac{H}{\log H}\right)$$
and t large of size X^2.
3. This is a consequence of the Vinogradov-Korobov zero-free region when $H > \exp((\log X)^{2/3+\varepsilon})$.

The "local" step

I will now explain the ideas in the "local" step. We are trying to show that if

$$\Big|\sum_{n\in I}\lambda(n)e(n\alpha_I)\Big| \gg H$$

for a positive proportion of intervals I, then, for a positive proportion of I,

$$\Big|\sum_{n\in I}\lambda(n)e(n\alpha_I)\Big| \approx \Big|\sum_{n\in I}\lambda(n)n^{iT_I}\Big|$$

and moreover

1. We have $|T_I| \ll X^2$
2. If $\frac{I}{p} \cap \frac{J}{q} \neq \emptyset$ for $p, q \in [P, 2P]$ then $T_I = T_J$.

The local step : The main tools

Our two main tools are :

1. (The large sieve) If

$$\sup_{J \cap I \neq \emptyset} \Big| \sum_{n \in J} \lambda(n) e(n\alpha) \Big| \geq \eta H \tag{5}$$

   then α is within $O(1/H)$ of a set of at most $O_\eta(1)$ elements. This morally allows us to assume that (up to perturbation by $O(1/H)$) and modulo 1) α is uniquely determined, and $\alpha_J = \alpha_I$ for any J such that $J \cap I \neq \emptyset$.

2. (Turan-Kubilius) For "most" primes p

$$\Big| \sum_{n \in I} \lambda(n) e(n\alpha_I) \Big| \approx p \Big| \sum_{n \in I/p} \lambda(np) e(np\alpha_I) \Big| = p \Big| \sum_{n \in I/p} \lambda(n) e(np\alpha_I) \Big|.$$

   Let us pretend this holds for all primes $p \in [P, 2P]$.

Consequently if I and J are such that $\frac{I}{p} \cap \frac{J}{q} \neq \emptyset$ then $p\alpha_I \equiv q\alpha_J + O(P/H) \pmod 1$. This is the main relation that we will use.

The "local step"

1. The fact that we only know $p\alpha_I \equiv q\alpha_J + O(P/H)$ modulo 1 is very restrictive.
2. It would be great if we could divide by q as having $(p/q)\alpha_I \equiv \alpha_J + O(1/H) \pmod 1$ is equivalent to $p\alpha_I \equiv q\alpha_J + O(P/H) \pmod q$ and thus contains more information.
3. One natural idea is that we can shift the α_J's in whatever way we wish by integers. Perhaps we can show that there exists $\widetilde{\alpha}_J = \alpha_J + k_J$ with $k_J \in \mathbb{Z}$ such that

$$\frac{p}{q}\widetilde{\alpha}_I = \widetilde{\alpha}_J + O(1/H) \pmod 1$$

   whenever $\frac{I}{p} \cap \frac{J}{q} \neq \emptyset$ and $p, q \in [P, 2P]$.
4. We were not quite able to prove that (only something weaker). Nonetheless let us pretend that we have the above relationship as it allows to succintly give the flavor of the rest of the proof, eliminating a lot of somewhat overwhelming technicalities but keeping the overall structure.

The "local" step

Let's assume that we have $\frac{p}{q}\alpha_I \equiv \alpha_J + O(1/H) \pmod 1$ whenever $\frac{I}{p} \cap \frac{J}{q} \neq \emptyset$ for $p, q \in [P, 2P]$. Pick $k \in 2\mathbb{N}$ such that $P^k \approx X$.

1. An application of Hölder (actually Sidorenko's conjecture) shows that if k is large then there are many k-cycles in the graph connecting I, J when $\frac{I}{p} \cap \frac{J}{q} \neq \emptyset$.
2. Iterating along this cycle we obtain (essentially) that $\alpha_I \ll X$ so we can write $\alpha_I = \frac{T_I}{x_I}$ with x_I the starting point of the interval I and $T_I \ll X^2$. In particular $e(n\alpha_I) \approx e^{i\theta_I} n^{iT_I}$ for $n \in I$.
3. This is not enough because it gives no relationship between distinct I and J.
4. We actually use the fact that there are many subgraphs that correspond to two k-cycles connected by an edge. The connecting edge is what allows us to show that $T_I = T_J$ whenever $\frac{I}{p} \cap \frac{J}{q} \neq \emptyset$ for primes $p, q \in [P, 2P]$.

Conclusion

1. The local step shows that if

$$\Big|\sum_{n\in I}\lambda(n)e(n\alpha_I)\Big| \gg H$$

then there exists a positive proportion subsets of intervals such that whenever I, J lie in it we have,

$$e(n\alpha_I) \approx e^{i\theta_I} n^{iT_I} \text{ for } n \in I \text{ and } e(n\alpha_J) \approx e^{i\theta_J} n^{iT_J} \text{ for } n \in J$$

and $T_I = T_J$ whenever $\frac{I}{p} \cap \frac{J}{q} \neq \emptyset$.

2. The global step shows that interpreting this as a graph property there exists a $T \ll X^2$ such that for a positive proportion of I we have $T_I = T$.
3. We conclude that for a positive proportion of intervals I,

$$\Big|\sum_{n\in I}\lambda(n)e(n\alpha_I)\Big| \approx \Big|\sum_{n\in I}\lambda(n)n^{iT}\Big|$$

4. This is contradicted by my result with Matomäki.

Thank you

Thank you!